

## **Horseshoe Priors for Edge-Preserving Linear Bayesian Inversion**

Uribe Felipe, Dong Yiqiu, Hansen Per Christian

This is a Author's accepted manuscript (AAM) version of a publication  
published by Society for Industrial and Applied Mathematics  
in SIAM Journal on Scientific Computing

**DOI:** DOI number of the original publication

### **Copyright of the original publication:**

© 2023 Society for Industrial and Applied Mathematics

### **Please cite the publication as follows:**

Uribe, F., Dong, Y., Hansen, P. C. (2023). Horseshoe Priors for Edge-Preserving Linear Bayesian Inversion. SIAM Journal on Scientific Computing, vol. 45, iss. 3. DOI: 10.1137/22M151036

# Horseshoe priors for edge-preserving linear Bayesian inversion

Felipe Uribe<sup>a,\*</sup>, Yiqiu Dong<sup>b</sup>, Per Christian Hansen<sup>b</sup>

<sup>a</sup>*School of Engineering Science, Lappeenranta-Lahti University of Technology. Yliopistonkatu 34, 53850 Lappeenranta, Finland.*

<sup>b</sup>*Department of Applied Mathematics and Computer Science, Technical University of Denmark. Richard Petersens Plads, Building 324, 2800 Kgs. Lyngby, Denmark.*

---

## Abstract

In many large scale inverse problems, such as computed tomography and image deblurring, characterization of sharp edges in the solution is desired. Within the Bayesian approach to inverse problems, edge-preservation is often achieved using Markov random field priors based on heavy-tailed distributions. Another strategy, popular in sparse statistical modeling, is the application of hierarchical shrinkage priors. An advantage of this formulation lies in expressing the prior as a conditionally Gaussian distribution depending on *global* and *local* hyperparameters which are endowed with heavy-tailed hyperpriors. In this work, we revisit the shrinkage horseshoe prior and introduce its formulation for edge-preserving settings. We discuss a Gibbs sampling framework to solve the resulting hierarchical formulation of the Bayesian inverse problem. In particular, one of the conditional distributions is high-dimensional Gaussian, and the rest are derived in closed form by using a scale mixture representation of the heavy-tailed hyperpriors. Applications from imaging science show that our computational procedure is able to compute sharp edge-preserving posterior point estimates with reduced uncertainty.

**Keywords:** Bayesian inverse problems, Bayesian hierarchical modeling, edge-preserving estimation, horseshoe prior, Gibbs sampler.

**MSC:** 62F15, 65C05, 65R32, 65F22.

---

## 1. Introduction

The integration of data into computational models is fundamental when one is interested in recovering unknown parameters defining the model. This specifies an inverse problem where the goal is to discover the parameters that make the model closely match the observed data. Inverse problems are typically ill-posed and the stability mainly depends on the structure of the mathematical model, the dimension of the parameter space, and the scarcity and noisiness of the data. Regularization methods are commonly used to find the solution of inverse problems (see, e.g., [1]). These methods are deterministic and incorporate penalty functions as stabilization procedure. Another approach relies on Bayesian statistics (see, e.g., [2]). These techniques incorporate a probabilistic description of the model parameters that combines prior information with a likelihood function that accounts for the model and data. In this approach, the objective is to estimate a so-called posterior distribution of the model parameters. Closed-form expressions of the posterior can only be derived in some particular cases, and in general the posterior has to be estimated using sampling-based methods such as Markov chain Monte Carlo (MCMC) [3, 4], or approximation methods such as variational inference [5], transport maps [6] and Laplace approximations [7, 8].

The main advantage of the Bayesian formulation lies in the possibility of quantifying parameter uncertainty due to noise and model errors. However, the approach is generally limited to problems where the dimension of the parameter space is not prohibitively large. In large-scale scenarios, the unknown parameters are modeled as random fields discretized pointwise on a fine grid. This complicates the application of sampling methods for Bayesian inference since one requires the exploration of a high-dimensional parameter space—even in the case of the linear problems considered here. Another level of complexity arises when the solution of the inverse problem is non-smooth and the preservation

---

\*Corresponding author.

Email address: felipe.uribe@lut.fi (Felipe Uribe)

of the edges or sharp features in the solution is necessary [9]. This especially occurs in imaging science such as X-ray computed tomography, image deblurring, segmentation and denoising. In these cases edge-preservation is generally imposed via the prior probability distribution. These prior models can be grouped into four categories:

1. *Heavy-tailed Markov random fields* defined on pairwise parameter increments or differences. The idea is to increase the probability of large jump events by imposing heavy-tailed distributions on the increments. Some examples include total variation prior [10], Laplace Markov random fields [11, 8] and Cauchy Markov random fields [12].
2. *Random fields with jump discontinuities*, including Besov space priors based on Haar wavelet expansions [13, 14], Gaussian and compound Poisson process priors [15], and level-set priors that employ Gaussian random fields with different thresholds via predefined level sets [16].
3. *Machine learning-based models*, including implicit models such as plug-and-play priors [17] and Bayesian neural networks with Cauchy weights that promote edge-preservation [18].
4. *Shrinkage priors*, also known as global-local or component-wise priors, that aim at shrinking small values towards zero while leaving the larger ones unaffected. These prior models are hierarchical by nature and include the methods proposed in [19, 20], as well as multiple models used in the statistics community such as elastic net, spike-slab, horseshoe, discrete Gaussian mixtures, and others (see, e.g., [21, 22] for reviews).

One of the advantages of shrinkage priors is that they are essentially defined by conditionally Gaussian distributions. Given one or more hyperparameters such as the variance, a target uncertain parameter endowed with a Gaussian distribution can have many nice properties [23]. Different shrinkage priors differ in the definition of the hierarchical structure and probabilistic models used to represent the hyperparameters. For example, the Lasso prior is based on a scale mixture representation of the Laplace distribution, and thus it defines an exponential distribution on the variance with a half-Cauchy distributed rate parameter [24]. The elastic net prior uses a truncated-gamma distribution on the variance with scale defined by two additional parameters that are half-Cauchy distributed [25]. Another shrinkage model is the horseshoe prior, in which the variance is defined by two hyperparameters that are half-Cauchy distributed [26]; one hyperparameter is a scalar that controls the global variability, while the other is spatially dependent and captures local variations.

Despite being common in statistics, shrinkage priors are not yet fully applied in large-scale linear inverse problems that require edge-preserving solutions. In this paper, we formulate a Bayesian inference approach that targets this requirement. We focus on the horseshoe prior, where the target parameter is Gaussian-distributed conditioned on half-Cauchy-distributed global and local hyperparameters. Our idea is that when this type of prior is formulated for the parameter increments or differences, the local hyperparameter assists in the identification of large discontinuities. The Bayesian inverse problem becomes hierarchical, since hyperparameters associated with the prior and likelihood are also part of the inference process.

The main difficulty in this formulation is that the resulting hierarchical posterior is challenging to handle computationally due to the large-scale nature of the target parameter and the heavy-tailed distributions imposed on the hyperparameters. To alleviate the complexity of the sampling process, we use the scale mixture representation of the half-Student-t distribution to express the half-Cauchy distributions of the hyperparameters in terms of inverse gamma distributions. Despite this requires the addition of extra hyperparameters, the hierarchical posterior is given essentially by products of Gaussian and inverse gamma densities. Therefore, we can exploit the conjugacy in the hierarchies to define a Gibbs sampler for which the associated full conditional densities are derived in closed form, and hence they can be sampled directly without an MCMC algorithm. However, within the Gibbs sampler, the conditional distribution of the target parameter defines a linear Gaussian Bayesian inverse problem. Although this full conditional is Gaussian and has an analytical expression, it is prohibitively slow to simulate in large-scale applications. Hence, the task of sampling from this distribution is formulated as a least-squares problem that is solved efficiently via standard iterative methods such as CGLS [27] with a preconditioner based on the prior precision matrix. The computational framework is tested on a one-dimensional deconvolution problem, as well as, two-dimensional applications of image deblurring and computed tomography (CT).

*Related work.* Existing applications of shrinkage priors in the context of inverse problems include the following. Bayesian inverse problems in sparse signal and image processing are surveyed in [28] where a variety of shrinkage priors are discussed and applied to multiple problems in imaging science. MCMC and approximation methods are also suggested in [28] to solve the Bayesian inverse problem depending on the

choice of the prior. Spike-and-slab priors are used in [29] for linear Bayesian inverse problems subject to sparsity constraints. The posterior is simulated with an expectation propagation algorithm using different approximations of the precision matrix of the underlying Gaussian parameter to improve the efficiency in high-dimensional applications. So-called *Gaussian hypermodels* are proposed in [19], which are defined as conditionally Gaussian priors with a single local variance hyperparameter endowed with a gamma distribution. An iterative alternating sequential algorithm is proposed that converges to the maximum a posteriori probability estimator of the parameters. Following this idea, in [20] a generalized gamma distribution is imposed on the local variances and propose two modifications to their iterative algorithm to exploit the global convexity ensured by gamma hyperpriors and the stronger sparsity promotion of the generalized gamma hyperpriors.

*Our contributions.* The main contributions of the paper are highlighted as follows:

- (i) We review the horseshoe prior for the representation of parameters that are sparse.
- (ii) We introduce the horseshoe prior for the solution of linear Bayesian inverse problems that require edge-preservation and perform uncertainty quantification.
- (iii) We propose a Gibbs sampler to solve the resulting hierarchical formulation of the Bayesian inverse problem; the full conditional distributions are derived in closed form by exploiting conjugacy.
- (iv) We demonstrate the performance of our method through numerical experiments on two inverse problems.

*Structure of the paper.* This paper is organized as follows. In section 2, we briefly introduce the mathematical framework for linear inverse problems and the Bayesian approach, and we present the standard horseshoe prior. We then derive its generalization for edge-preserving linear Bayesian inversion in section 3. We show that a scale mixture representation can be applied to express a heavy-tailed half-Cauchy hyperprior in terms of two inverse gamma distributions. The computational framework is defined in section 4, where we develop a Gibbs sampling approach to solve the resulting hierarchical Bayesian inverse problem. Here, we use CGLS to sample the full Gaussian conditional and we derive the remaining hyperparameter conditionals in closed form. In section 5 we illustrate our computational approach with imaging science applications, before concluding in section 6.

## 2. Preliminaries

We commence with the mathematical framework and introduce linear inverse problems and the Bayesian approach to find solutions to those, as well as the standard horseshoe prior.

### 2.1. Linear inverse problems

We consider the discrete inverse problem of estimating an unknown model parameter  $\mathbf{x} \in \mathbb{R}^d$  using noisy observed data  $\mathbf{y} \in \mathbb{R}^m$ , given a linear forward operator  $\mathbf{A} \in \mathbb{R}^{m \times d}$  acting as a link between the parameters and data. Here,  $d$  is the number of parameters and  $m$  denotes the number of data points. The parameter  $\mathbf{x}$  is spatially varying and exists in connection with a physical domain  $D$ , e.g., one-dimensional signals, two-dimensional images and three-dimensional objects. Using a suitable finite-dimensional representation, the dimension  $d$  will depend on a discretization size  $N$  imposed on the physical domain. For instance, in one- and two-dimensional settings  $d = N$  and  $d = N^2$  (assuming equal discretization in both directions), respectively.

In particular, a linear inverse problem estimates an  $\mathbf{x}$  that approximates the ground truth  $\mathbf{x}^{\text{true}}$  in

$$\mathbf{y} = \mathbf{A}\mathbf{x}^{\text{true}} + \mathbf{e} \quad \text{with} \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \sigma_{\text{obs}}^2 \mathbf{I}_m), \quad (1)$$

where the additive measurement noise is modeled as a realization of an independent Gaussian random vector with variance  $\sigma_{\text{obs}}^2$ ,  $\mathbf{I}_m$  is an identity matrix of size  $m \times m$ , and  $\mathcal{N}(\cdot, \cdot)$  denotes the multivariate Gaussian distribution depending on mean and covariance parameters. The linear inverse problem associated to (1) is typically ill-posed and we employ the Bayesian statistical framework to study the influence of noise in the observed data.

## 2.2. Bayesian inverse problems

In the Bayesian approach to inverse problems [2], we treat  $\mathbf{x}$ ,  $\mathbf{y}$  and  $\mathbf{e}$  as random variables and the solution is expressed as the probability distribution of  $\mathbf{x}$  given an instance of the observed data  $\mathbf{y}$ . This allows both modeling the noise via its statistical properties and specifying prior information on the parameter, i.e., the form of solutions that are believed more likely.

The unknown parameter is modeled as a discretized random field represented as a random vector  $\mathbf{X}$  taking values  $\mathbf{x} \in \mathbb{R}^d$ . We assume the distribution of  $\mathbf{X}$  has a so-called *prior* probability density  $\pi_{\text{pr}}(\mathbf{x})$ . The *likelihood function* is defined from a density  $\pi_{\text{data}}(\cdot | \mathbf{x})$  with fixed argument equal to the observed data  $\mathbf{y}$ ; then the likelihood becomes a function of  $\mathbf{x}$  only. The Gaussian statistical model for the measurement noise assumed in (1) gives

$$\pi_{\text{data}}(\mathbf{y} | \mathbf{x}) = \frac{1}{(2\pi)^{m/2} \sigma_{\text{obs}}^m} \exp\left(-\frac{1}{2\sigma_{\text{obs}}^2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2\right). \quad (2)$$

Assuming that the measurement noise and model parameters are independent, the solution of the Bayesian inverse problem combines the prior and likelihood probability models into a so-called *posterior* density, given by Bayes' Theorem as

$$\pi_{\text{pos}}(\mathbf{x} | \mathbf{y}) = \frac{1}{Z} \pi_{\text{data}}(\mathbf{y} | \mathbf{x}) \pi_{\text{pr}}(\mathbf{x}), \quad (3)$$

where  $Z = \int_{\mathbb{R}^d} \pi_{\text{data}}(\mathbf{y} | \mathbf{x}) \pi_{\text{pr}}(\mathbf{x}) d\mathbf{x}$  is the normalizing constant of the posterior.

In practice, point estimates of the posterior (3) are used to represent the solution of the inverse problem. The two classical choices are the *maximum a posteriori* (MAP) and the *posterior mean* (PM) estimators:

$$\mathbf{x}_{\text{MAP}} = \arg \max_{\mathbf{x} \in \mathbb{R}^d} \pi_{\text{pos}}(\mathbf{x} | \mathbf{y}) \quad \text{and} \quad \mathbf{x}_{\text{PM}} = \int_{\mathbb{R}^d} \mathbf{x} \pi_{\text{pos}}(\mathbf{x} | \mathbf{y}) d\mathbf{x}. \quad (4)$$

The MAP estimator is computed via optimization techniques which is often an advantage, compared to the calculation of the PM estimator that requires a more involved high-dimensional integration. The estimation of the posterior mean and related summary statistics is often performed via Monte Carlo methods (see, e.g., [3, 4]). We point out that when the first two statistical moments of a random vector do not exist (as it is common for some heavy-tailed distribution models), location and scale characteristics of the distribution can still be summarized using for instance the *posterior median* and *median absolute deviation*, respectively [30]. These are also a sensible option to describe random variables whose distribution is heavily skewed.

To find the posterior (3), we will focus on the definition of a prior probability model that allows both, the preservation of potential sharp features in the unknown parameter and a tractable computation of the high-dimensional posterior distribution. Our model is based on the horseshoe prior [31], which we revisit next.

## 2.3. Horseshoe prior

When our *a priori* knowledge is that  $\mathbf{x}$  is sparse, it may be convenient to apply a sparsity-inducing shrinkage prior (see, e.g., [22]). A well-known continuous shrinkage model is the standard horseshoe [31], which is defined as the hierarchical prior:

$$\pi_{\text{pr}}(\mathbf{x}, \tau, \boldsymbol{\sigma}) = \pi_{\text{pr}}(\mathbf{x} | \tau, \boldsymbol{\sigma}) \pi_{\text{hpr}}(\tau) \pi_{\text{hpr}}(\boldsymbol{\sigma}), \quad (5)$$

where a zero-mean conditionally Gaussian prior is imposed on  $\mathbf{x}$ :

$$\pi_{\text{pr}}(\mathbf{x} | \tau, \boldsymbol{\sigma}) = \frac{1}{(2\pi)^{d/2} (\det \boldsymbol{\Sigma}(\tau, \boldsymbol{\sigma}))^{1/2}} \exp\left(-\frac{1}{2} \mathbf{x}^T \boldsymbol{\Sigma}^{-1}(\tau, \boldsymbol{\sigma}) \mathbf{x}\right), \quad \boldsymbol{\Sigma}(\tau, \boldsymbol{\sigma}) = \tau^2 \text{diag}(\sigma_1^2, \dots, \sigma_d^2). \quad (6)$$

Here, the prior covariance matrix  $\boldsymbol{\Sigma}(\tau, \boldsymbol{\sigma}) \in \mathbb{R}^{d \times d}$  depends on the hyperparameters  $\tau \in \mathbb{R}_{>0}$  controlling *global shrinkage* and  $\boldsymbol{\sigma} = [\sigma_1, \dots, \sigma_d]^T \in \mathbb{R}_{>0}^d$  defining *local shrinkage*. The global-local scheme is defined with heavy-tailed distributions. In particular, the horseshoe model uses a half-Cauchy distribution for  $\tau$  and an independent standard half-Cauchy distribution for  $\boldsymbol{\sigma}$ :

$$\pi_{\text{hpr}}(\tau) \propto \frac{2}{\tau_0 \left(1 + \frac{\tau^2}{\tau_0^2}\right)} \quad \text{and} \quad \pi_{\text{hpr}}(\boldsymbol{\sigma}) \propto \prod_{i=1}^d \frac{2}{1 + \sigma_i^2} \quad \text{with } \tau, \sigma_i > 0, \quad (7)$$

where  $\tau_0$  is a scale parameter. According to existing work, we can set (i)  $\tau_0 \approx \sigma_{\text{obs}}$  [31]; (ii)  $\tau_0 = (d_0/(d - d_0))\sigma_{\text{obs}}$ , where  $d_0$  is the number of nonzero elements [32]; (iii)  $\tau_0$  as a part of the inference, e.g., using a Jeffreys' hyperprior [26]. In this paper, we numerically compare the influence of the choice of  $\tau_0$ .

Another expression for the horseshoe prior is given by re-writing the density of  $\sigma_i$  as  $\pi_{\text{hpr}}(\sigma_i) \propto 1/(1 + \sigma_i^2) = \kappa_i$  for  $i = 1, \dots, d$ , where  $\kappa_i$  is called a shrinkage parameter. The relation between  $\sigma_i$  and  $\kappa_i$  is bijective when  $\sigma_i > 0$ , then its inverse is continuously differentiable and non-zero. Therefore, the hyperprior of each parameter  $\kappa_i$  is found from the half-Cauchy hyperprior of  $\sigma_i$  using a standard probabilistic transformation and we obtain

$$\pi_{\text{hpr}}(\kappa_i) \propto \frac{1}{\sqrt{\kappa_i} \sqrt{1 - \kappa_i}}, \quad (8)$$

which is proportional to a *horseshoe-shaped* beta probability density with shape parameters equal to  $1/2$ , hence the name of the prior. Small values of  $\kappa_i$  correspond to  $\sigma_i \rightarrow \infty$  and produce almost no shrinkage, while values  $\kappa_i \approx 1$  corresponding to  $\sigma_i \rightarrow 0$  provide essentially full shrinkage, and hence  $\kappa_i$  can be used to describe how many active or inactive variables are present in the model [26].

The horseshoe prior is hierarchical by definition and lacks of a proper analytical expression after marginalizing out the hyperparameters. However, lower and upper bounds for the one-dimensional horseshoe probability density are available [26]:

$$\frac{1}{2\sqrt{2\pi^3}} \log \left( 1 + \frac{4}{x^2} \right) \leq \pi_{\text{pr}}(x) \leq \frac{1}{\sqrt{2\pi^3}} \log \left( 1 + \frac{2}{x^2} \right) \quad \text{for } x \neq 0. \quad (9)$$

These bounds are shown in Figure 1 together with other common probability distributions (standardized). Note that the horseshoe prior has heavy tails that are comparable to the Cauchy distribution, and it has a pole at  $x = 0$ . These characteristics typically enable the prior to perform well when handling sparsity. We exploit this behavior to develop a horseshoe-based model that can be applied to a class of linear Bayesian inverse problems where edge-preservation is fundamental.




Figure 1: Comparison of the horseshoe prior bounds (shaded area) with other common probability densities. The zoom-in highlights the distributions at the tails.

### 3. Horseshoe prior for edge-preservation

In some applications the solution of inverse problems is sparse; in other applications sharp edges are present in the solution meaning that pairwise differences or increments of the solution elements exhibit sparsity. We exploit this fact to define a horseshoe prior on the differences that promote solutions preserving their sharp features. The resulting prior can be interpreted as an anisotropic conditionally Gaussian Markov random field prior.

#### 3.1. Horseshoe prior and conditionally Gaussian Markov random fields

Gaussian Markov random fields (GMRF) are typically specified by their conditional independence structure. Therefore, a natural way to describe a GMRF is by its precision matrix. This is because the sparsity structure of this matrix determines the so-called neighborhood system explaining conditional relations in  $\mathbf{x}$  [33]. Since we are interested in preserving the edges, we define the precision matrix of the GMRF using increments between elements of the parameter vector. In one-dimensional cases, the

increments are computed by application of a finite difference matrix  $\mathbf{D}$ , while in two-dimensional cases we consider the horizontal and vertical first-order finite difference matrices  $\mathbf{D}^{(1)}, \mathbf{D}^{(2)}$ . These are given by

$$\mathbf{D} = \begin{bmatrix} 1 & & & & \\ -1 & 1 & & & \\ & -1 & 1 & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \end{bmatrix}_{N \times N} \quad \text{and} \quad \mathbf{D}^{(1)} = \mathbf{I}_N \otimes \mathbf{D}, \quad \mathbf{D}^{(2)} = \mathbf{D} \otimes \mathbf{I}_N, \quad (10)$$

where  $N$  is domain discretization size,  $\mathbf{I}_N$  is the identity matrix of size  $N \times N$ , and  $\otimes$  denotes the Kronecker product. Note that we use, without loss of generality, zero boundary conditions on  $\mathbf{x}$  at the left part of the domain.

The structure defined in (10) considers pairwise differences (i.e., in 1D the vector  $\mathbf{D}\mathbf{x}$ , and in 2D the vectors  $\mathbf{D}^{(1)}\mathbf{x}$  and  $\mathbf{D}^{(2)}\mathbf{x}$ ). We assume that the elements in these random vectors are independent. The idea now is to impose a horseshoe prior on them. This allows us to rewrite the prior in (6) as a conditionally GMRF that incorporates information about the increments. The associated probability density is

$$\pi_{\text{pr}}(\mathbf{x} \mid \tau, \mathbf{w}) = \frac{(\det \mathbf{\Lambda}(\tau, \mathbf{w}))^{1/2}}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2}\mathbf{x}^\top \mathbf{\Lambda}(\tau, \mathbf{w}) \mathbf{x}\right), \quad (11)$$

where  $\mathbf{\Lambda}(\tau, \mathbf{w})$  is a prior precision matrix (see below) that depends on the global standard deviation  $\tau \in \mathbb{R}_{>0}$  and local weights  $\mathbf{w} = [w_1, \dots, w_d]^\top \in \mathbb{R}_{>0}^d$ . These hyperparameters are defined as in the standard horseshoe prior and they are half-Cauchy-distributed following (7). In this case, the weights in  $\mathbf{w}$  are analogous to local standard deviations in  $\boldsymbol{\sigma}$ , but they will correspond to the increments and not to the elements of the parameter vector, hence the different notation.

For the conditionally Gaussian prior in (11), the precision matrix is defined such that it takes into account the increments via (10) (see, e.g., [34, p.68] for the details). Therefore, we can generally write the precision matrix as

$$\mathbf{\Lambda}(\tau, \mathbf{w}) = \mathbf{L}^\top \mathbf{W}(\tau, \mathbf{w}) \mathbf{L}, \quad (12)$$

and for one-dimensional and two-dimensional domains we have

$$\mathbf{L} = \mathbf{D}, \quad \mathbf{W}(\tau, \mathbf{w}) = \begin{bmatrix} \frac{1}{\tau^2 w_1^2} & & \\ & \ddots & \\ & & \frac{1}{\tau^2 w_d^2} \end{bmatrix} \quad (13)$$

and

$$\mathbf{L} = \begin{bmatrix} \mathbf{D}^{(1)} \\ \mathbf{D}^{(2)} \end{bmatrix}, \quad \mathbf{W}(\tau, \mathbf{w}) = \begin{bmatrix} \mathbf{W}^{(1)} & \mathbf{0}_d \\ \mathbf{0}_d & \mathbf{W}^{(2)} \end{bmatrix}, \quad \mathbf{W}^{(i)} = \begin{bmatrix} \frac{1}{\tau^2 (w_1^{(i)})^2} & & \\ & \ddots & \\ & & \frac{1}{\tau^2 (w_d^{(i)})^2} \end{bmatrix}, \quad i = 1, 2, \quad (14)$$

respectively. Here,  $\mathbf{L}$  is called the structure matrix, and  $\mathbf{0}_d \in \mathbb{R}^{d \times d}$  is a matrix containing zeros. Notice that in the two-dimensional case we have to consider local weights

$$\mathbf{w} = \begin{bmatrix} \mathbf{w}^{(1)} \\ \mathbf{w}^{(2)} \end{bmatrix} \in \mathbb{R}_{>0}^{2d}$$

for each coordinate direction. These are assumed independent and identically distributed according to the half-Cauchy prior exactly as in the one-dimensional case. We remark that for GMRFs the (prior) precision matrix (12) is sparse, which can be exploited in numerical computations.

The conditionally Gaussian prior (11), based on precision matrices of the form (12) with (13)–(14), and the hyperpriors (7) define a horseshoe prior promoting edge-preservation. Again, this is because we are now regularizing *increments* of the unknown parameter using a heavy-tailed probabilistic model. A main motivation to utilize the horseshoe prior is that the unknown parameter is conditionally Gaussian, which facilitates the analytical and numerical treatment of the resulting posterior distribution. However, the main difficulties in applying this prior are: (i) the dimension of the parameter space is increased since the local weights have the same or even larger dimension than the model parameter, and (ii) the hyperparameters are endowed with heavy-tailed distributions that cause challenges when computing



the posterior via sampling methods. For the latter point, a *regularized horseshoe prior* is proposed in [32]; this model employs half-Student's t-distributions with larger degrees of freedom on the local parameters to overcome the sampling issues. However, increasing the degrees of freedom makes the tail of the distribution less heavy, which is not desirable when looking for methods that promote sharp features.

### 3.2. Extended horseshoe prior

Our objective is to present an equivalent model that extends the hierarchical structure of the horseshoe prior by adding auxiliary parameters. This allows us to write the resulting posterior such that it can be sampled in a tractable manner without losing the connection to the original half-Cauchy hyperpriors in the standard horseshoe. We point out that this idea is also used in [35] for the standard horseshoe prior applied to logistic regression problems.

We define a scale mixture decomposition of a Student's t-distributed random variable (see, e.g., [36]). If  $A$  and  $B$  are random variables such that

$$(A^2 | B) \sim \text{IG}\left(\frac{\nu}{2}, \frac{\nu}{B}\right) \quad \text{and} \quad B \sim \text{IG}\left(\frac{1}{2}, \frac{1}{c^2}\right), \quad \text{then} \quad A \sim t^+(\nu, 0, c), \quad (15)$$

where  $\text{IG}(\cdot, \cdot)$  denotes the inverse gamma distribution with shape and scale parameters, and  $t^+(\cdot, \cdot, \cdot)$  is the half-Student's t-distribution depending on degrees of freedom, location and scale parameters. Both densities are defined in appendix Appendix A, along with a simple derivation of the mixture representation given in (15).

The scale mixture representation (15) is utilized to write the standard horseshoe prior as an extended hierarchical prior. The motivation behind this choice is that the additional auxiliary variables allow the derivation of full conditional distributions in closed form by exploiting conjugacy. Thereafter, these can be sampled within a Gibbs sampler scheme [37, Ch.10]. The resulting *extended horseshoe prior* is defined as:

$$\pi_{\text{pr}}(\mathbf{x}, \tau^2, \mathbf{w}^2, \gamma, \boldsymbol{\xi}) = \pi_{\text{pr}}(\mathbf{x} | \tau^2, \mathbf{w}^2) \pi_{\text{hpr}}(\tau^2 | \gamma) \pi_{\text{hpr}}(\gamma) \pi_{\text{hpr}}(\mathbf{w}^2 | \boldsymbol{\xi}) \pi_{\text{hpr}}(\boldsymbol{\xi}), \quad (16)$$

where

$$(\mathbf{x} | \tau^2, \mathbf{w}^2) \sim \mathcal{N}(\mathbf{0}, \mathbf{\Lambda}^{-1}(\tau, \mathbf{w})), \quad (17a)$$

$$(\tau^2 | \gamma) \sim \text{IG}\left(\frac{\nu}{2}, \frac{\nu}{\gamma}\right), \quad \gamma \sim \text{IG}\left(\frac{1}{2}, \frac{1}{\tau_0^2}\right), \quad (17b)$$

$$(w_i^2 | \xi_i) \sim \text{IG}\left(\frac{\nu}{2}, \frac{\nu}{\xi_i}\right), \quad \xi_i \sim \text{IG}\left(\frac{1}{2}, 1\right), \quad (17c)$$

with the squaring  $\mathbf{w}^2$  taken elementwise.

We will focus on the standard horseshoe prior with degrees of freedom  $\nu = 1$ , such that the half-Student's t-distribution becomes a half-Cauchy distribution. However, we note that the regularized horseshoe prior [32], which typically uses  $\nu = 3$ , can also be included in the formulation.

## 4. Proposed approach

The extended hierarchical horseshoe prior (16) is used to define the posterior in (3), thereby defining a hierarchical formulation of the Bayesian inverse problem. In addition to the hyperparameters arising from the horseshoe prior, we also model the noise variance  $\sigma_{\text{obs}}^2$  in the likelihood function as a random variable. Following the hyperprior distributions imposed on the prior parameters, we define an inverse gamma hyperprior  $\pi_{\text{hpr}}(\sigma_{\text{obs}}^2)$  with shape and scale parameters  $\alpha_{\text{obs}} = 1$  and  $\beta_{\text{obs}} = 10^{-4}$ , respectively. This choice makes the hyperprior relatively uninformative [38].

As a result, the Bayesian inverse problem of estimating the posterior (3) is written as the hierarchical Bayesian inverse problem of determining the new posterior density:

$$\pi_{\text{pos}}(\mathbf{x}, \sigma_{\text{obs}}^2, \tau^2, \mathbf{w}^2, \gamma, \boldsymbol{\xi} | \mathbf{y}) \propto \pi_{\text{data}}(\mathbf{y} | \mathbf{x}, \sigma_{\text{obs}}^2) \pi_{\text{pr}}(\mathbf{x} | \tau^2, \mathbf{w}^2) \pi_{\text{hpr}}(\tau^2 | \gamma) \pi_{\text{hpr}}(\mathbf{w}^2 | \boldsymbol{\xi}) \pi_{\text{hpr}}(\sigma_{\text{obs}}^2) \pi_{\text{hpr}}(\gamma) \pi_{\text{hpr}}(\boldsymbol{\xi}); \quad (18)$$

whose dependencies are shown in Figure 2.

Performing statistical inference with the posterior (18) generally requires application of sampling methods based on Markov chain Monte Carlo (MCMC) methods. The goal in MCMC is to compute samples or realizations of a Markov chain that is stationary with respect to the posterior distribution [4].






Figure 2: Structure of the hierarchical Bayesian inference model (18) based on the extended horseshoe prior. The nodes reflect the uncertain parameters and their dimension. Starting from the data  $\mathbf{y}$ , the solid edges reflect dependencies in the likelihood (on  $\mathbf{x}$  and  $\sigma_{\text{obs}}^2$ ) and prior distribution (on the hyperparameters  $\mathbf{w}^2, \tau^2, \boldsymbol{\xi}, \gamma$ ). Here, the dimension of the local parameter vector is  $k = d$  or  $k = 2d$ , if the problem is one- or two-dimensional, respectively.

However, hierarchical Bayesian inverse problems commonly lead to posterior densities that resemble a funnel for which small variations of the hyperparameters induce large changes in the posterior density values. This is a standard difficulty for traditional MCMC algorithms (see, e.g., [39] for a detailed discussion). To alleviate this problem one typically relies on (i) re-parametrization of the hierarchical structure in order to break potential correlations between parameters (cf. Figure 2), or (ii) application of specialized MCMC algorithms that capture the geometry of the joint posterior. For instance, a Riemannian Hamiltonian Monte Carlo method that uses local curvature information is advocated in [39]; and for problems involving the horseshoe prior, a Gibbs sampler is proposed in [35] and two further MCMC methods are developed in [40]. Recently, another Gibbs sampler is discussed in [41] for Bayesian inverse problems under the regularized horseshoe prior.

Our idea is to exploit the structure presented in section 3.2 to derive full conditional distributions for each uncertain parameter. The advantage of this approach is that it allows a direct application of the Gibbs sampler [42], since the conditional densities for each parameter can be derived in closed form. Therefore, we avoid sampling in the high-dimensional joint space of the posterior. In particular, the full conditional densities associated with (18) are found to be (cf. fig. 2):

$$\pi_1(\mathbf{x} \mid \sigma_{\text{obs}}^2, \tau^2, \mathbf{w}^2) \propto \pi_{\text{data}}(\mathbf{y} \mid \mathbf{x}, \sigma_{\text{obs}}^2) \pi_{\text{pr}}(\mathbf{x} \mid \tau^2, \mathbf{w}^2), \quad (19a)$$

$$\pi_2(\sigma_{\text{obs}}^2 \mid \mathbf{x}) \propto \pi_{\text{data}}(\mathbf{y} \mid \mathbf{x}, \sigma_{\text{obs}}^2) \pi_{\text{hpr}}(\sigma_{\text{obs}}^2), \quad (19b)$$

$$\pi_3(\tau^2 \mid \mathbf{x}, \mathbf{w}^2, \gamma) \propto \pi_{\text{pr}}(\mathbf{x} \mid \tau^2, \mathbf{w}^2) \pi_{\text{hpr}}(\tau^2 \mid \gamma), \quad (19c)$$

$$\pi_4(\mathbf{w}^2 \mid \mathbf{x}, \tau^2, \boldsymbol{\xi}) \propto \pi_{\text{pr}}(\mathbf{x} \mid \tau^2, \mathbf{w}^2) \pi_{\text{hpr}}(\mathbf{w}^2 \mid \boldsymbol{\xi}), \quad (19d)$$

$$\pi_5(\gamma \mid \tau^2) \propto \pi_{\text{hpr}}(\tau^2 \mid \gamma) \pi_{\text{hpr}}(\gamma), \quad (19e)$$

$$\pi_6(\boldsymbol{\xi} \mid \mathbf{w}^2) \propto \pi_{\text{hpr}}(\mathbf{w}^2 \mid \boldsymbol{\xi}) \pi_{\text{hpr}}(\boldsymbol{\xi}). \quad (19f)$$

In the remainder of this section, we determine the full conditional densities for each uncertain parameter and discuss sampling techniques to obtain draws from them. Due to the extended horseshoe model, we anticipate that most of the densities in (19) can be sampled directly and a Gibbs sampler can be used to characterize the posterior.

#### 4.1. Sampling of $\pi_1$

The conditional for the uncertain parameter  $\mathbf{x}$  in (19a) defines a Bayesian inverse problem with Gaussian likelihood and prior. After inserting the corresponding probabilistic models we obtain:

$$\pi_1(\mathbf{x} \mid \sigma_{\text{obs}}^2, \tau^2, \mathbf{w}^2) \propto \exp \left( -\frac{1}{2} \left( \frac{1}{\sigma_{\text{obs}}^2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \|\boldsymbol{\Lambda}^{1/2}(\tau, \mathbf{w}) \mathbf{x}\|_2^2 \right) \right). \quad (20)$$

This conditional density is also Gaussian with precision matrix and mean vector given by (see, e.g., [2, p.78])

$$\tilde{\boldsymbol{\Lambda}}(\tau, \mathbf{w}) = \frac{1}{\sigma_{\text{obs}}^2} \mathbf{A}^\top \mathbf{A} + \boldsymbol{\Lambda}(\tau, \mathbf{w}), \quad \tilde{\boldsymbol{\mu}}(\tau, \mathbf{w}) = \tilde{\boldsymbol{\Lambda}}^{-1}(\tau, \mathbf{w}) \left( \frac{1}{\sigma_{\text{obs}}^2} \mathbf{A}^\top \mathbf{y} \right), \quad (21)$$

where  $\boldsymbol{\Lambda}(\tau, \mathbf{w})$  is given in (12).

Given the hyperparameters  $\sigma_{\text{obs}}$ ,  $\tau$  and  $\mathbf{w}$ , the most direct sampling algorithm for a Gaussian distribution is based on the Cholesky factorization. In this case, a sample from  $\pi_1$  is obtained as  $\mathbf{x}^* = \tilde{\boldsymbol{\mu}} + \tilde{\boldsymbol{\Lambda}}^{-1/2} \mathbf{u}$ , where  $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$  is a standard Gaussian random vector, and  $\tilde{\boldsymbol{\Lambda}}^{1/2}$  is a lower triangular matrix with real and positive diagonal entries (Cholesky factor). Note that we drop the

dependence on  $\sigma_{\text{obs}}$ ,  $\tau$  and  $\mathbf{w}$  for the sake of notation. We refer to [43, 23] for a review of methods for sampling high-dimensional Gaussian distributions. The Cholesky strategy is computationally prohibitive due to the need for computing a factorization of the matrix  $\tilde{\mathbf{\Lambda}}(\tau, \mathbf{w})$ . Assuming that matrix-vector multiplications with  $\tilde{\mathbf{\Lambda}}$  can be done efficiently, we rely on Krylov subspace methods to sample from the Gaussian conditional in (20). In particular, the task of sampling a Gaussian random vector can be written as a least squares problem, and thus, we draw a sample  $\mathbf{x}^*$  from  $\pi_1$  by solving:

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{R}^d} \|\mathbf{M}\mathbf{x} - \mathbf{z}\|_2^2 \quad \text{with} \quad \mathbf{M} = \begin{bmatrix} (1/\sigma_{\text{obs}})\mathbf{A} \\ \mathbf{\Lambda}^{1/2} \end{bmatrix}, \quad \mathbf{z} = \begin{bmatrix} (1/\sigma_{\text{obs}})\mathbf{y} \\ \mathbf{0}_d \end{bmatrix} + \tilde{\mathbf{u}}, \quad (22)$$

where  $\tilde{\mathbf{u}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{m+d})$ . Recall that the solution of (22) is required at every Gibbs iteration given new values of the hyperparameters. We use the CGLS method, which requires one pair of forward and backward model computations (i.e., multiplications with  $\mathbf{A}$  and  $\mathbf{A}^\top$ ) per iteration. A tolerance  $\varepsilon_{\text{cgl}}s$  on the relative residual norm of the normal equations or a maximum number of iterations  $n_{\text{max}}$  can be used to control the quality of the computed samples.

One way to accelerate the performance of the CGLS method is to apply a standard-form transformation to (22) [1, Sec. 8.4]; this is also referred to as *priorconditioning* [44]. The idea is to compute the Cholesky factorization of the prior precision matrix  $\mathbf{\Lambda} = \mathbf{C}^\top \mathbf{C}$  and introduce the change of variables  $\tilde{\mathbf{x}} = \mathbf{C}\mathbf{x}$ , such that we can transform (22) into

$$\tilde{\mathbf{x}}^* = \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} \|\mathbf{M}\tilde{\mathbf{x}} - \mathbf{z}\|_2^2 \quad \text{with} \quad \mathbf{M} = \begin{bmatrix} (1/\sigma_{\text{obs}})\mathbf{A}\mathbf{C}^{-1} \\ \mathbf{I}_d \end{bmatrix}, \quad (23)$$

where  $\mathbf{z}$  is defined as in (22). Note that the change of variables corresponds to “whitening” the prior, that is,  $\tilde{\mathbf{x}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ . In the standard-form least squares problem (23), we are now working with the Krylov subspace  $\text{span}\{\mathbf{\Lambda}^{-1}\mathbf{A}^\top\mathbf{y}, (\mathbf{\Lambda}^{-1}\mathbf{A}^\top\mathbf{A})\mathbf{\Lambda}^{-1}\mathbf{A}^\top\mathbf{y}, (\mathbf{\Lambda}^{-1}\mathbf{A}^\top\mathbf{A})^2\mathbf{\Lambda}^{-1}\mathbf{A}^\top\mathbf{y}, \dots\}$ , which is in most cases a better subspace than  $\text{span}\{\mathbf{A}^\top\mathbf{y}, (\mathbf{A}^\top\mathbf{A})\mathbf{A}^\top\mathbf{y}, (\mathbf{A}^\top\mathbf{A})^2\mathbf{A}^\top\mathbf{y}, \dots\}$  of the general-form least-squares problem (22) (see, e.g., [1, Sec. 8.4] for the details). Hence, the application of the standard-form transformation tends to reduce the number of iterations required for convergence and it can be incorporated in a preconditioned version of the CGLS algorithm ([27, Sec. 7.4]). We denote this variant as pCGLS.

The Cholesky factor of the prior precision matrix, which is instrumental for the standard-form transformation, can be obtained as follows. In the one-dimensional case, from the definition of the prior precision in (13) we have that

$$\mathbf{\Lambda} = \mathbf{L}^\top \mathbf{W} \mathbf{L} = \mathbf{C}^\top \mathbf{C} \quad \text{with} \quad \mathbf{L} = \mathbf{D}, \quad (24)$$

where  $\mathbf{W}$  is diagonal and  $\mathbf{C}$  is either upper or lower triangular. Hence, by simple observation, we can write  $\mathbf{C} = \mathbf{W}^{1/2} \mathbf{D}$ , which is indeed triangular and constitutes an inexpensive representation of the Cholesky factor.

In the two-dimensional case, the Cholesky factor is no longer immediately available but still can be computed in an economical way. From (14) we have

$$\mathbf{\Lambda} = \mathbf{L}^\top \mathbf{W} \mathbf{L} \quad \text{with} \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}^{(1)} & \mathbf{0}_d \\ \mathbf{0}_d & \mathbf{W}^{(2)} \end{bmatrix} \quad \text{and} \quad \mathbf{L} = \begin{bmatrix} \mathbf{D}^{(1)} \\ \mathbf{D}^{(2)} \end{bmatrix}, \quad (25)$$

where the difference matrices  $\mathbf{D}^{(1)}$  and  $\mathbf{D}^{(2)}$  are sparse and so is  $\mathbf{L}$ . Now, we can compute the thin QR factorization of  $\mathbf{W}^{1/2} \mathbf{L}$ ,

$$\mathbf{W}^{1/2} \mathbf{L} = \begin{bmatrix} (\mathbf{W}^{(1)})^{1/2} \mathbf{D}^{(1)} \\ (\mathbf{W}^{(2)})^{1/2} \mathbf{D}^{(2)} \end{bmatrix} = \mathbf{Q} \mathbf{R}; \quad (26)$$

notice that  $\mathbf{Q}$  is not required and  $\mathbf{L}^\top \mathbf{W} \mathbf{L} = \mathbf{R}^\top \mathbf{R}$ . Therefore,  $\mathbf{C} = \mathbf{R}$  is the desired Cholesky factor.

We remark that despite the sparse QR factorization simplifies the presentations of the algorithm, it does not scale with the dimension. Different preconditioners can be used here for sampling Gaussian distributions in high dimensions, e.g., [43].

#### 4.2. Sampling of $\pi_2$ to $\pi_6$

The conditional densities for each hyperparameter are all obtained in closed form. This is due to conjugate relations that arise from the extended horseshoe prior formulation.

The conditional density for the *noise variance* in (19b) is

$$\pi_2(\sigma_{\text{obs}}^2 \mid \mathbf{x}) \propto \pi_{\text{data}}(\mathbf{y} \mid \mathbf{x}, \sigma_{\text{obs}}^2) \pi_{\text{hpr}}(\sigma_{\text{obs}}^2) \quad (27a)$$

$$\propto (\sigma_{\text{obs}}^2)^{-m/2} \exp\left(-\frac{1}{2\sigma_{\text{obs}}^2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2\right) \left[(\sigma_{\text{obs}}^2)^{-\alpha_{\text{obs}}-1} \exp\left(-\frac{\beta_{\text{obs}}}{\sigma_{\text{obs}}^2}\right)\right] \quad (27b)$$

$$\propto (\sigma_{\text{obs}}^2)^{-m/2-\alpha_{\text{obs}}-1} \exp\left(-\frac{1}{\sigma_{\text{obs}}^2} \left[\frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \beta_{\text{obs}}\right]\right), \quad (27c)$$

$$(\sigma_{\text{obs}}^2 | \mathbf{x}) \sim \text{IG}\left(\frac{m}{2} + \alpha_{\text{obs}}, \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \beta_{\text{obs}}\right). \quad (27d)$$

The conditional density for the squared *global shrinkage parameter* in (19c) is

$$\pi_3(\tau^2 | \mathbf{x}, \mathbf{w}^2, \gamma) \propto \pi_{\text{pr}}(\mathbf{x} | \tau^2, \mathbf{w}^2) \pi_{\text{hpr}}(\tau^2 | \gamma) \quad (28a)$$

$$\propto (\tau^2 \mathbf{w}^2)^{-k/2} \exp\left(-\frac{1}{2} \sum_{i=1}^k \frac{[\mathbf{L}\mathbf{x}]_i^2}{\tau^2 w_i^2}\right) \left[(\tau^2)^{-\nu/2-1} \exp\left(-\frac{\nu}{\gamma \tau^2}\right)\right] \quad (28b)$$

$$\propto (\tau^2)^{-k/2-\nu/2-1} \exp\left(-\frac{1}{\tau^2} \left[\frac{1}{2} \sum_{i=1}^k \frac{[\mathbf{L}\mathbf{x}]_i^2}{w_i^2} + \frac{\nu}{\gamma}\right]\right), \quad (28c)$$

$$(\tau^2 | \mathbf{x}, \mathbf{w}^2, \gamma) \sim \text{IG}\left(\frac{k+\nu}{2}, \frac{1}{2} \sum_{i=1}^k \frac{[\mathbf{L}\mathbf{x}]_i^2}{w_i^2} + \frac{\nu}{\gamma}\right), \quad (28d)$$

where  $k = d$  or  $k = 2d$  in one- and two-dimensional problems, respectively.

Moreover, since the *local shrinkage parameters* are assumed independent, one can derive the conditional density for their squared version and at each component (19d), as follows

$$\pi_4(w_i^2 | \mathbf{x}, \tau^2, \xi_i) \propto \pi_{\text{pr}}(\mathbf{x} | \tau^2, w_i^2) \pi_{\text{hpr}}(w_i^2 | \xi_i) \quad (29a)$$

$$\propto (\tau^2, w_i^2)^{-1/2} \exp\left(-\frac{1}{2} \frac{[\mathbf{L}\mathbf{x}]_i^2}{\tau^2 w_i^2}\right) \left[(w_i^2)^{-\nu/2-1} \exp\left(-\frac{\nu}{\xi_i w_i^2}\right)\right] \quad (29b)$$

$$\propto (w_i^2)^{-1/2-\nu/2-1} \exp\left(-\frac{1}{w_i^2} \left[\frac{[\mathbf{L}\mathbf{x}]_i^2}{2\tau^2} + \frac{\nu}{\xi_i}\right]\right), \quad (29c)$$

$$(w_i^2 | \mathbf{x}, \tau^2, \xi_i) \sim \text{IG}\left(\frac{\nu+1}{2}, \frac{[\mathbf{L}\mathbf{x}]_i^2}{2\tau^2} + \frac{\nu}{\xi_i}\right). \quad (29d)$$

For the *auxiliary parameters*, the conditionals  $\pi_5(\gamma | \tau^2)$  and  $\pi_6(\xi_i | w_i^2)$  in (19e) and (19f), respectively, are derived in a similar manner. These are inverse gamma distributions such that:

$$(\gamma | \tau^2) \sim \text{IG}\left(\frac{\nu+1}{2}, \frac{1}{\tau_0^2} + \frac{\nu}{\tau^2}\right), \quad (30a)$$

$$(\xi_i | w_i^2) \sim \text{IG}\left(\frac{\nu+1}{2}, 1 + \frac{\nu}{w_i^2}\right). \quad (30b)$$

#### 4.3. The computational procedure

Based on the sampling approaches for the full conditional densities discussed above, we define a Gibbs sampler generating  $n_s$  states of a Markov chain with stationary distribution (18),  $\{\mathbf{x}^{(j)}, (\sigma_{\text{obs}}^2)^{(j)}, (\tau^2)^{(j)}, (\mathbf{w}^2)^{(j)}, \gamma^{(j)}, \boldsymbol{\xi}^{(j)}\}_{j=1}^{n_s}$ .

The Gibbs sampler draws a single sample from each conditional density at each iteration. This uses the fact that, under mild conditions, the set of full conditional distributions determine the joint distribution [45]. The Markov chain approaches its equilibrium condition as the number of iterations increases. This means that after convergence all samples from the chain will be distributed according to the target posterior distribution. Convergence conditions for the Gibbs sampler are defined in [46, 47]. From a practical viewpoint, with a finite number of samples, we require application of burn-in and lagging steps: (i) the  $n_s$  successive values of the chain are only selected after discarding  $n_b$  samples during the warm up phase of the algorithm (burn-in period), and (ii) the sample chain values are only stored every  $n_t$  iterations since for large enough  $n_t$  the samples are virtually independent (lagging steps). As a result, to obtain  $n_s$  quasi-independent samples from the posterior, we require  $n = n_b + n_s n_t$  Markov chain steps.

Different scanning strategies exist for the Gibbs sampler (see, e.g., [3, 4]). We follow a deterministic or systematic scan in which all iterations consist of sampling the conditional densities of each component in the same order. This version of the Gibbs sampler applied to the posterior (18) is summarized in Algorithm 1. We remark that geometric ergodicity for different Gibbs samplers in the context of the

---

**Algorithm 1: GIBBS SAMPLING OF THE POSTERIOR (18)**

---

**Input:** conditional densities (19), number of samples  $n_s$ , thinning  $n_t$ , burn-in  $n_b$ , and maximum number of CGLS iterations  $n_{\max}$  (or tolerance  $\varepsilon_{\text{cglS}}$ ).

```
1 Initial states  $\mathbf{x}^*, \sigma_{\text{obs}}^*, \tau^*, \mathbf{w}^*, \gamma^*, \boldsymbol{\xi}^*$ , and initial Cholesky factor  $\mathbf{C}(\tau^*, \mathbf{w}^*)$ .
2  $n \leftarrow n_b + n_s n_t$ , and  $j \leftarrow 1$ 
3 for  $k = 1, \dots, n$  do
4   // Sample target parameter
5    $\mathbf{x}^* \sim \pi_1(\cdot \mid \sigma_{\text{obs}}^*, \tau^*, \mathbf{w}^*)$ : solve least-squares problem using preconditioned CGLS with  $\mathbf{C}(\tau^*, \mathbf{w}^*)$ 
6   // Sample variance hyperparameters
7    $\sigma_{\text{obs}}^* \sim \pi_2(\cdot \mid \mathbf{x}^*)$ : in closed form (27)
8    $\tau^* \sim \pi_3(\cdot \mid \mathbf{x}^*, \mathbf{w}^*, \gamma^*)$ : in closed form (28d)
9    $\mathbf{w}^* \sim \pi_4(\cdot \mid \mathbf{x}^*, \tau^*, \boldsymbol{\xi}^*)$ : in closed form (29d)
10  // Compute Cholesky factor
11   $\mathbf{C}(\tau^*, \mathbf{w}^*)$ : based on the approach discussed in section 4.1
12  // Sample auxiliary hyperparameters
13   $\gamma^* \sim \pi_5(\cdot \mid \tau^*)$ : in closed form (30a)
14   $\boldsymbol{\xi}^* \sim \pi_6(\cdot \mid \mathbf{w}^*)$ : in closed form (30b)
15  // Save samples
16  if  $(k > n_b)$  then
17    if  $(k \bmod n_t) = 0$  then
18       $\mathbf{x}^{(j)} \leftarrow \mathbf{x}^*, (\sigma_{\text{obs}}^2)^{(j)} \leftarrow \sigma_{\text{obs}}^*, (\tau^2)^{(j)} \leftarrow \tau^*, (\mathbf{w}^2)^{(j)} \leftarrow \mathbf{w}^*, \gamma^{(j)} \leftarrow \gamma^*, \boldsymbol{\xi}^{(j)} \leftarrow \boldsymbol{\xi}^*$ 
19       $j \leftarrow j + 1$ 
20    end
21  end
22 end
23 return  $\{\mathbf{x}^{(j)}, (\sigma_{\text{obs}}^2)^{(j)}, (\tau^2)^{(j)}, (\mathbf{w}^2)^{(j)}, \gamma^{(j)}, \boldsymbol{\xi}^{(j)}\}_{j=1}^{n_s}$ 
```

---

horseshoe prior has been shown in [48, 41]; this provides indications on the accuracy of asymptotic standard errors in the MCMC-based posterior estimates.

Furthermore, the efficiency of Algorithm 1 can be measured in terms of the autocorrelation of the posterior samples. The chain autocorrelation allows the definition of the *effective sample size* [4]

$$n_{\text{eff}} := \frac{n_s}{1 + 2 \sum_{j=1}^{n_s} \frac{\rho^{(j)}}{\rho^{(0)}}} \approx \left\lceil \frac{n_s}{\tau_{\text{int}}} \right\rceil, \quad (31)$$

where  $\rho^{(j)}$  denotes the autocorrelation at the  $j$ th lag and  $\tau_{\text{int}}$  is the *integrated autocorrelation time* (IACT). Essentially, the effective sample size is used to compare the variance estimated via correlated MCMC samples and the ideal case of a variance computed from independent draws. Thus, the aim is to obtain a value of  $n_{\text{eff}}$  as close as possible to  $n_s$ .

The computational cost of the Gibbs sampler at each iteration can be given in terms of the number of model calls (operations with  $\mathbf{A}$  and  $\mathbf{A}^T$ ). Hence, the cost corresponding to the simulation of each conditional density are: maximum  $2n_{\max}$  calls are needed for  $\pi_1$  and 1 evaluation is required to sample the density  $\pi_2$ . Hence, the total cost of a single iteration of the Gibbs sampler is  $2n_{\max} + 1$  model calls maximum.

## 5. Numerical experiments

In the following, we illustrate the use of the horseshoe prior for edge-preserving inversion and the Gibbs sampler in Algorithm 1 designed for the computations. We consider linear inverse problems arising in imaging science. The first example consists of a one-dimensional deconvolution problem that allows us to test multiple parameter settings in our computational framework. The second example is a two-dimensional computed tomography problem that highlights the potential of the horseshoe prior in more realistic applications.

In our test problems, the point estimates  $\bar{\mathbf{x}}$  (posterior mean or median) of the target parameter are evaluated using the relative reconstruction error defined as

$$\text{relerr} := \|\bar{\mathbf{x}} - \mathbf{x}^{\text{true}}\|_2 / \|\mathbf{x}^{\text{true}}\|_2, \quad (32)$$

where  $\mathbf{x}^{\text{true}}$  denotes the underlying true solution.

To further assess the effectiveness of our approach, we compare our solutions with those from the method proposed in [8] which uses a Laplace Markov random field prior to achieve edge-preservation. This method, however, relies on a Laplace approximation of the posterior, while Algorithm 1 does not introduce any approximation.

### 5.1. One-dimensional deconvolution

We consider the inverse problem of identifying an unknown piecewise constant signal  $x : [0, 1] \rightarrow \mathbb{R}$  from noisy convolved data. The mathematical model for convolution can be written as a Fredholm integral equation of the first kind:

$$y(t) = \int_0^1 \mathcal{A}(t-u)x(u) du \quad \text{with} \quad \mathcal{A}(t) = \exp\left(-\frac{1}{2s^2}(t-u)^2\right), \quad 0 \leq t \leq 1, \quad (33)$$

where  $y(t)$  denotes the convolved signal and we employ a Gaussian convolution kernel  $\mathcal{A}(t)$  with fixed parameter  $s = 0.016$ .

In practice, a finite-dimensional representation of (33) is employed. After discretizing the signal domain into  $d = 128$  intervals, the convolution model can be expressed as a system of linear algebraic equations  $\mathbf{y} = \mathbf{A}\mathbf{x}$ . We consider two sets of synthetic observed data  $\mathbf{y}$  with  $m = d$  equally-spaced elements. The first set is generated with noise standard deviation  $\sigma_{\text{obs}}^{\text{true}} = 7.867 \times 10^{-3}$  and the second one with  $\sigma_{\text{obs}}^{\text{true}} = 1.967 \times 10^{-2}$  (corresponding to 2% and 5% errors). Figure 3 shows the true signal together with the two data sets.




Figure 3: Deconvolution example. Left: underlying true signal. Center: convolved and noisy data (dots) with 2% noise. Right: convolved and noisy data with 5% noise.

The inputs to the Gibbs sampler are studied/selected as follows: (i) The number of samples is  $n_s = 2 \times 10^4$ , in addition to a burn-in period of  $n_b = 2 \times 10^3$ . (ii) The scale parameter  $\tau_0$  in the half-Cauchy hyperprior for  $\tau$  is studied. (iii) The lag or thinning number  $n_t$  is also analyzed. (iv) The influence of the number of standard and preconditioned CGLS iterations is studied as well.

We remark that for this example we can afford the factorization of the precision matrix  $\tilde{\mathbf{A}}$  in (21). Hence, we can directly sample the Gaussian conditional  $\pi_1$  in (20) without using CGLS. This allows us to obtain an exact sample from  $\pi_1$  at each Gibbs iteration. This method is referred to as “direct” in the following studies and will be used to compare to the solution obtained by sampling  $\pi_1$  with standard and preconditioned CGLS.

*Influence of the scale parameter  $\tau_0$ .* In this case, we use the low-noise data set, fix the lagging steps to  $n_t = 20$ , and we directly sample the conditional  $\pi_1$ . Moreover, we fix the noise standard deviation to its true value. We perform a parameter study on  $\tau_0$  using the values  $\{(d_0/(d-d_0))\sigma_{\text{obs}}^{\text{true}}, \sigma_{\text{obs}}^{\text{true}}, 1\}$ , where we assume the number of relevant variables is  $d_0 = 10$ , and thus,  $\tau_0 \in \{6.667 \times 10^{-4}, 7.867 \times 10^{-3}, 1\}$ . The Gibbs sampler is applied for each value and we analyze the output of the Markov chain associated with the global parameter  $\tau$ , as well as the relative reconstruction error. For the point estimators, we also report the median since we observe that the posterior distribution of the  $\tau$  parameter is right-skewed. The median is also used for the parameter  $\mathbf{x}$  as it generates smaller reconstruction errors than the mean estimator.

The results are shown in Table 1. We observe that for the smallest value of  $\tau_0$ , the best reconstruction errors are obtained, however, the effective sample size  $n_{\text{eff}}$  decreases. There are slight variations in the posterior mean and median of  $\tau$  but overall the value remains in the same order of magnitude. In general, for this particular example, we do not observe significant changes in the results by modifying the parameter  $\tau_0$ . Nevertheless, from this simple study, we notice that the standard recommendation  $\tau_0 \approx \sigma_{\text{obs}}^{\text{true}}$  covers a middle ground in terms of the effective sample size and reconstruction error values.

Table 1: Posterior mean and median of  $\tau$ , effective sampling size of the  $\tau$ -chain, and relative reconstruction error based on the mean and median estimators, for different values of the scale parameter  $\tau_0$ .

$\tau_0$	Mean	Median	$n_{\text{eff}}$	relerr (mean)	relerr (median)
$(d_0/(d - d_0))\sigma_{\text{obs}}^{\text{true}}$	$6.218 \times 10^{-3}$	$5.675 \times 10^{-3}$	4036	$1.413 \times 10^{-2}$	$1.099 \times 10^{-2}$
$\sigma_{\text{obs}}^{\text{true}}$	$7.603 \times 10^{-3}$	$7.059 \times 10^{-3}$	4896	$1.571 \times 10^{-2}$	$1.216 \times 10^{-2}$
1	$9.096 \times 10^{-3}$	$8.364 \times 10^{-3}$	4426	$1.712 \times 10^{-2}$	$1.311 \times 10^{-2}$

**Remark 5.1.** From this study, we select the value of the scale parameter  $\tau_0$  as the value of the noise standard deviation  $\sigma_{\text{obs}}$ . Recall, however, that the true noise variance is unknown and hence we model it as a hyperparameter with associated conditional density  $\pi_2$  in (27). Therefore, the value of  $\tau_0^2$  is now equal to a realization of  $\sigma_{\text{obs}}^2$  at a given Gibbs iteration. Due to this modeling choice, the conditional density for the auxiliary parameter  $\gamma$  in (30a) will now depend on  $\sigma_{\text{obs}}^2$ , i.e.,  $\pi_5(\gamma \mid \tau^2, \tau_0^2 = \sigma_{\text{obs}}^2)$ , and hence, the sampling of  $\gamma$  in line 13 of Algorithm 1 can be modified accordingly.

*Influence of the thinning  $n_t$ .* As in the previous study, we use the small noise data set and we directly sample  $\pi_1$ . However, we now set the squared scale parameter  $\tau_0^2$  equal to realizations of the noise variance random variable (cf. Remark 5.1). This parameter study on the lag number  $n_t$  uses the values  $\{10, 20, 40, 80\}$ . Once again, the Gibbs sampler is implemented with each value and we analyze the output of the Markov chain associated to the one-dimensional hyperparameters  $\tau$  and  $\sigma_{\text{obs}}$ , and also the reconstruction error.

The results are shown in Table 2. The noise standard deviation  $\sigma_{\text{obs}}$  remains essentially unaffected by the thinning. For global parameter  $\tau$ , we observe that the IACT and consequently the number of effective samples improves with larger thinning. Particularly, we require a considerably high value of  $n_t$  (in this case 80 or more) to obtain quasi-independent samples from  $\tau$ . Nevertheless, note that the mean and standard deviation estimators remain mostly unchanged by the thinning. We also compute the relative reconstruction error based on the median estimator of  $\mathbf{x}$ , with increasing thinning. The values for each parameter  $n_t$  are  $\text{relerr} = \{1.213, 1.216, 1.215, 1.214\} \times 10^{-2}$ , thus, the errors are approximately the same independent of the choice of  $n_t$ . Based on these results and to strike a balance between the computational cost and the quality of the posterior samples, the lag number for the thinning step can be selected between 20 and 40.

Table 2: MCMC chain metrics and posterior statistics for the hyperparameters  $\sigma_{\text{obs}}$  and  $\tau$ , with increasing thinning values.

Hyperparam.	$n_t$	Mean	Std	IACT	$n_{\text{eff}}$
$\sigma_{\text{obs}}$	10	$7.945 \times 10^{-3}$	$5.585 \times 10^{-4}$	1.06	18 899
	20	$7.947 \times 10^{-3}$	$5.559 \times 10^{-4}$	1.03	19 507
	40	$7.949 \times 10^{-3}$	$5.584 \times 10^{-4}$	1.03	19 409
	80	$7.948 \times 10^{-3}$	$5.614 \times 10^{-4}$	1.05	19 036
$\tau$	10	$7.631 \times 10^{-3}$	$3.190 \times 10^{-3}$	8.13	2 460
	20	$7.609 \times 10^{-3}$	$3.168 \times 10^{-3}$	4.09	4 888
	40	$7.619 \times 10^{-3}$	$3.151 \times 10^{-3}$	2.23	8 986
	80	$7.596 \times 10^{-3}$	$3.145 \times 10^{-3}$	1.29	15 399

*Standard and preconditioned CGLS.* We use the settings obtained from the previous experiments, and now we select  $n_t = 40$ . We run the Gibbs sampler with the draws from  $\pi_1$  computed by standard and preconditioned CGLS. These are compared with the solution obtained by directly sampling from  $\pi_1$ . Both variants of CGLS are stopped at a given tolerance  $\varepsilon_{\text{cgl}} \in \{10^{-3}, 10^{-4}\}$ . To compare the influence of the tolerance value in CGLS, we compute the relative reconstruction error in the solution.

Figure 4 shows the evolution of the CGLS iterations  $n_{\text{cgl}}$  with the number of samples obtained at every Gibbs iteration. These are shown for each predefined tolerance in the CGLS algorithm. Lines track



the mean number of iterations, and for reference, the dotted vertical line marks the burn-in index. We directly observe that preconditioned CGLS requires considerably less number of iterations per sample, compared to standard CGLS.




Figure 4: Comparison of standard and preconditioned CGLS for different tolerances in the algorithm. The solid lines track the cumulative mean of the number of iterations. The dotted vertical line marks the burn-in period.

Table 3 shows a detailed comparison of standard and preconditioned CGLS using both tolerance values. The relative reconstruction errors for the mean and median estimates of  $\mathbf{x}$  are shown. To measure the quality of the results with respect to the posterior variability of  $\mathbf{x}$ , we compute the relative error of the standard deviation estimates; here we assume the standard deviation obtained by the direct method as true value. Moreover, the average number of CGLS/pCGLS iterations and the average effective sample size (across the  $d$  components of the  $\mathbf{x}$  chain) are also computed. We observe that with  $\varepsilon_{\text{cgl}} = 10^{-4}$  CGLS and pCGLS give comparable results, and naturally pCGLS requires much less iterations. However, when  $\varepsilon_{\text{cgl}} = 10^{-3}$ , the results of pCGLS deteriorate particularly in the estimation of the posterior variability. We see that the relative error in the posterior standard deviation is considerably large, while the errors in the mean and median estimates actually improve compared to the other cases. This indicates that, for this tolerance value, there is a concentration of the samples towards the mode, while missing the tails of the distribution. As a result, we require more stringent tolerances in pCGLS to guarantee convergence in order to obtain an exact sample from the Gaussian conditional  $\pi_1$ . In addition, we see that overall the samples of the parameter  $\mathbf{x}$  are almost independent, given the average  $n_{\text{eff}}$  values are close to  $n_s$ .

Table 3: Comparison of standard and preconditioned CGLS for different CGLS tolerances. The direct method that samples exactly from  $\pi_1$  is also shown.

		relerr (mean)	relerr (median)	relerr (std)	Average $n_{\text{cgl}}$	Average $n_{\text{eff}}$
	Direct	$1.56 \times 10^{-2}$	$1.21 \times 10^{-2}$	—	—	19386
$\varepsilon_{\text{cgl}} = 10^{-3}$	CGLS	$1.52 \times 10^{-2}$	$1.20 \times 10^{-2}$	$5.81 \times 10^{-2}$	143	19291
	pCGLS	$0.89 \times 10^{-2}$	$0.69 \times 10^{-2}$	$34.5 \times 10^{-2}$	33	18857
$\varepsilon_{\text{cgl}} = 10^{-4}$	CGLS	$1.55 \times 10^{-2}$	$1.21 \times 10^{-2}$	$1.99 \times 10^{-2}$	238	19384
	pCGLS	$1.54 \times 10^{-2}$	$1.19 \times 10^{-2}$	$1.78 \times 10^{-2}$	62	19369

*Solution for different noise levels.* Now we use both data sets. The solutions are computed using preconditioned CGLS with  $\varepsilon_{\text{cgl}} = 10^{-4}$ . First we show in Figure 5 posterior statistics for the hyperparameters  $\sigma_{\text{obs}}$  and  $\tau$  for the low noise data set. In particular, we plot the histogram of the posterior samples and generated chain, as well as, the ergodic mean and the sample autocorrelation. We observe that the posterior for the global parameter  $\tau$  is right-skewed, while the posterior of  $\sigma_{\text{obs}}$  is almost symmetric. The posterior samples computed by the Gibbs sampler have significantly small correlation and we also see that the posterior chain has converged and it is well-mixed. We obtain similar results for the large noise data set and thus we omit the figure.

Figure 6 shows the posterior mean and 95% credible interval for the local weight parameter  $\mathbf{w}$ . Note that the local parameter captures the locations where the signal edges are placed. The uncertainty is large at those particular locations and essentially zero in the rest of the domain. We observe that as the noise increases, the magnitude of the weights decreases and the point estimators become less sharp. Nevertheless, the posterior is still able to capture the edges correctly.



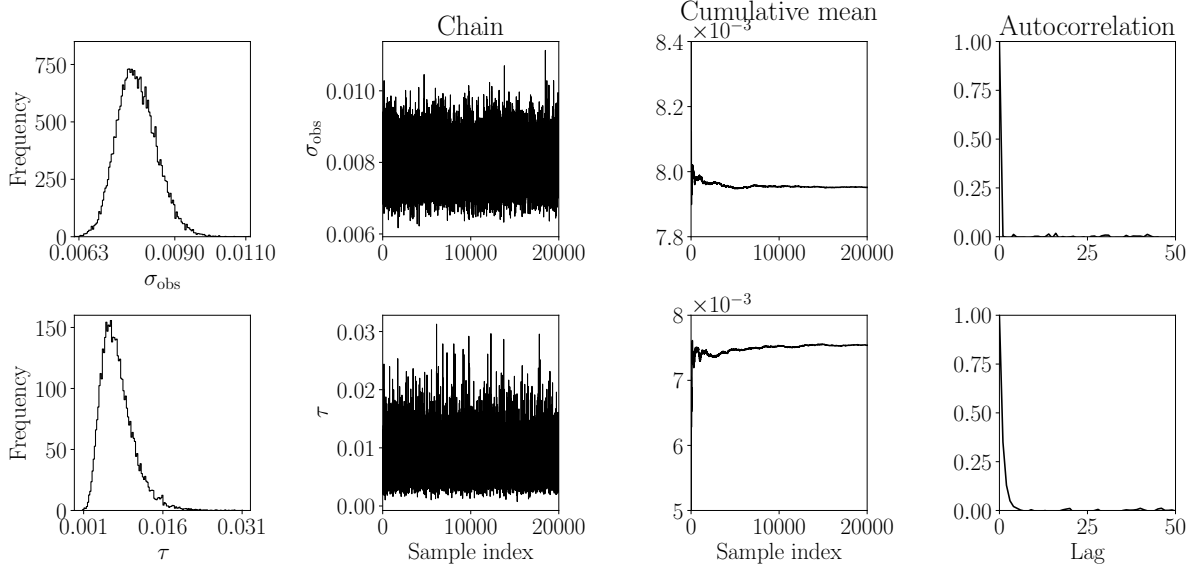


Figure 5: Posterior of  $\sigma_{\text{obs}}$  (first row) and  $\tau$  (second row) using the low noise data set. Histogram, chain, cumulative mean and sample autocorrelation.




Figure 6: Posterior statistics for  $w$  using the low noise (left) and high noise (right) data sets. Posterior mean (solid line) and 95% credible interval (shaded area).

Finally, the main posterior statistics for the target parameter  $x$  are shown in Figure 7 together with the underlying true signal. Similar to the local weights, we plot the posterior mean and 95% credible interval. The relative reconstruction errors based on the mean, for the low and high noise cases, are  $1.54 \times 10^{-2}$  and  $6.63 \times 10^{-2}$ , respectively. We compare the results based on the horseshoe prior to the method in [8] which is based on a Laplace Markov random field prior. For the latter, we use the same number of samples, burn-in and thinning. In the case of the Laplace Markov random field prior, we obtain the reconstruction errors  $5.36 \times 10^{-2}$  and  $9.27 \times 10^{-2}$ , for each data set. We see that our method not only generates sharper solutions but also reduces the posterior uncertainty.

## 5.2. Two-dimensional examples

In this section, we study a couple of linear inverse problems to test the horseshoe prior in two-dimensional applications. In both examples, the inputs to the Gibbs sampler are selected as follows: (i) The number of samples is  $n_s = 2 \times 10^4$ , in addition to a burn-in period of  $n_b = 4 \times 10^3$ . (ii) The scale parameter  $\tau_0$  is defined from the noise standard deviation as mentioned in Remark 5.1. (iii) The thinning number is set to  $n_t = 20$ . (iv) In two-dimensional applications the computation of the QR decomposition in (26) required in pCGLS dominates the computational cost. Instead, we design the experiments such that the direct method to sample  $\pi_1$  can be utilized, and for this we employ sparse matrix Cholesky factorizations using the `cholmod` package of the `scikit-sparse` library in python.

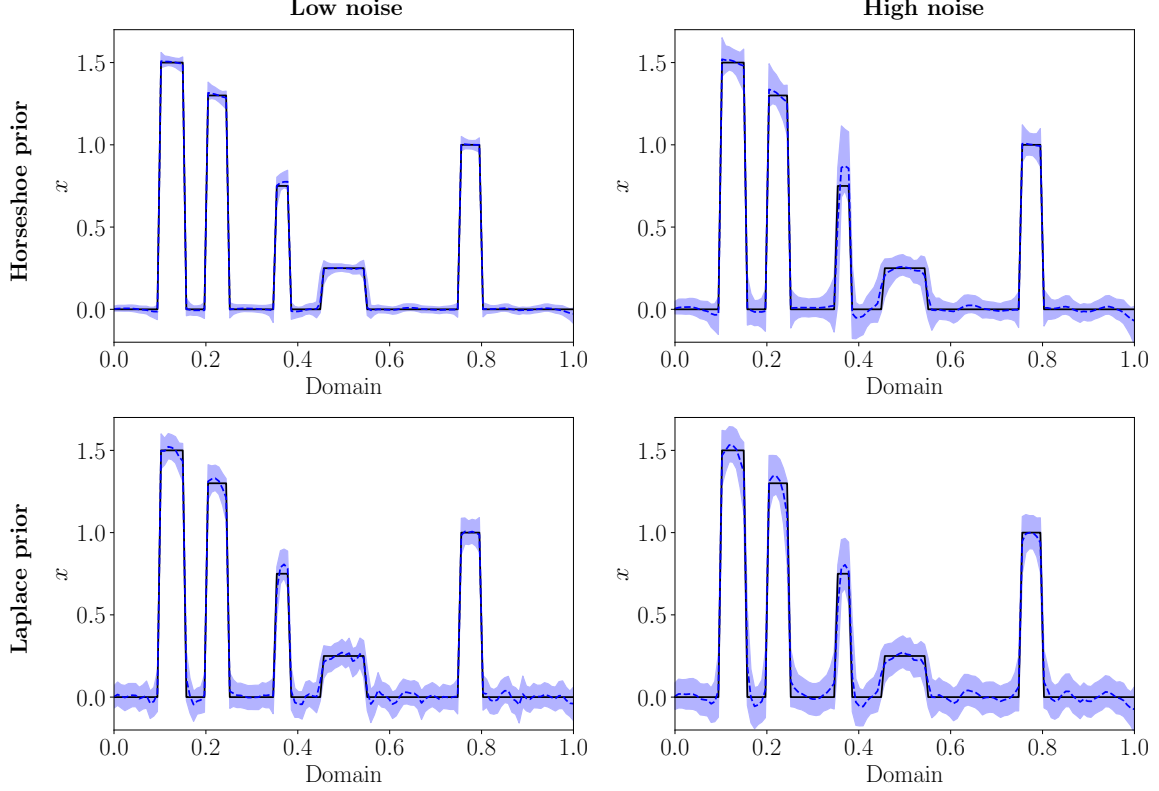


Figure 7: Comparison of posterior statistics for the target parameter  $\mathbf{x}$  using the horseshoe prior (first row) and the Laplace prior (second row). The results are shown using the low noise (left) and high noise (right) data sets. True solution (solid line), posterior mean (dashed line) and 95% credible interval (shaded area).

### 5.2.1. Image deblurring

We consider an image deblurring inverse problem where we seek to recover an original sharp image from a blurred one by using a mathematical model of the blurring process. This is analogous to one-dimensional deconvolution but now in two dimensions. In this case, we assume that the blurring of the columns in the image is independent of the blurring of the rows. Therefore, there exist two matrices  $\mathbf{A}_c \in \mathbb{R}^{N \times N}$  and  $\mathbf{A}_r \in \mathbb{R}^{N \times N}$ , such that we can express the forward operator as  $\mathbf{A}_c \mathbf{X} \mathbf{A}_r^T$  with an image  $\mathbf{X} \in \mathbb{R}^{N \times N}$ . Here, we set  $\mathbf{A}_c$  and  $\mathbf{A}_r$  as in Challenge 1 from [49]; in MATLAB notation these are:

```
c = zeros(N,1); c(1:5) = [5:-1:1]'/15; Ac = toeplitz(c)
r = zeros(N,1); r(1:10) = [5:-.5:.5]'/15; Ar = toeplitz(c,r).
```




Figure 8: Image deblurring example. Left: underlying true image. Right: blurred noisy image.

Since we typically work with a vectorized version of the image, we can define the system matrix as  $\mathbf{A} = (\mathbf{A}_r \otimes \mathbf{A}_c) \in \mathbb{R}^{d \times d}$ , such that the linear forward operator becomes  $\mathbf{A}\mathbf{x}$ . In this example, the image

size is 32-by-32, i.e.,  $N = 32$  and  $d = N^2 = 1024$ . Moreover, we set the Gaussian noise level to 1%. The underlying true image and its blurred and noisy version are shown in Figure 8.

The posterior solution for the hyperparameters  $\sigma_{\text{obs}}$  and  $\tau$  is shown in Figure 9. As in the previous example, we plot the histogram of the posterior samples, the Markov chain, the ergodic mean and the sample autocorrelation. The magnitude of both hyperparameters is very small, especially the global standard deviation  $\tau$  that tends to shrink parameters to zero. Furthermore, we observe that the posterior samples computed by the Gibbs sampler have small correlation and that the chain has converged and it is well-mixed.




Figure 9: Posterior of  $\sigma_{\text{obs}}$  (first row) and  $\tau$  (second row). Histogram, chain, cumulative mean and sample autocorrelation. Note the sharp decay in the autocorrelation after first lag, an indication of quasi-independent samples.

Figure 10 shows the posterior mean and standard deviation for the local weight parameter  $\mathbf{w}$ . Note that the local parameter captures the locations where the signal edges are placed. Particularly, the local parameter  $\mathbf{w}^{(1)}$  captures the locations where the image vertical edges are located, while the weight  $\mathbf{w}^{(2)}$  identifies the horizontal edges. Similar as in the one-dimensional example, the weights and their uncertainty are large at the edge locations in order to escape shrinkage, and essentially zero in the rest of the domain.

Posterior statistics for the image represented by  $\mathbf{x}$  are shown in Figure 11. We compare the posterior mean and standard deviation obtained by using the horseshoe prior with those from the Laplace Markov random field prior [8]. We show the posterior mean with the same intensity range in both results for the sake of comparison, however, there are no negative values in the horseshoe case. The relative reconstruction error based on the mean are 0.098 and 0.065, respectively. Despite the relative error is slightly larger for the horseshoe prior in this example, we clearly see that the mean of the posterior based on the horseshoe prior is sharper and has smaller posterior uncertainty. Nevertheless, with the horseshoe prior the standard deviation is particularly large at the locations that are not identified properly, while with the Laplace prior the larger values of standard deviation are located at the edges.

Finally, notice that for the posterior based on the horseshoe prior most of the error contribution is associated to the misrepresentation of a few corners of the objects in the image. To illustrate this we plot a few posterior samples in Figure 12. In some cases, the samples detect the corner pixels with interior intensity but others with the outside intensity (see, e.g., bottom right corner of the “round” object). We observe that with the horseshoe prior the edges are always sharp, but those arising from the Laplace prior are in general more noisy. This is related to the fact that the horseshoe prior is more heavy tailed than the Laplace prior.

### 5.2.2. X-ray CT reconstruction

The objective of 2D CT reconstruction is to estimate an image of the cross-section of an object using projection data that represents the intensity loss or attenuation of a beam of X-rays as they pass through the object.

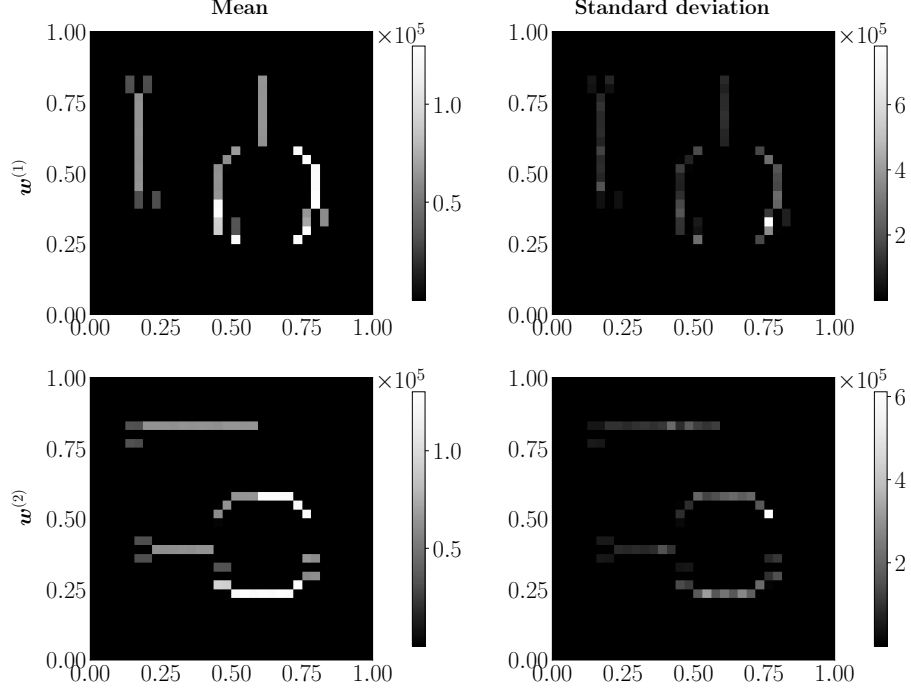


Figure 10: Posterior statistics for the local weights  $\mathbf{w}^{(1)}$  (first row) and  $\mathbf{w}^{(2)}$  (second row). Posterior mean (left) and standard deviation (right).

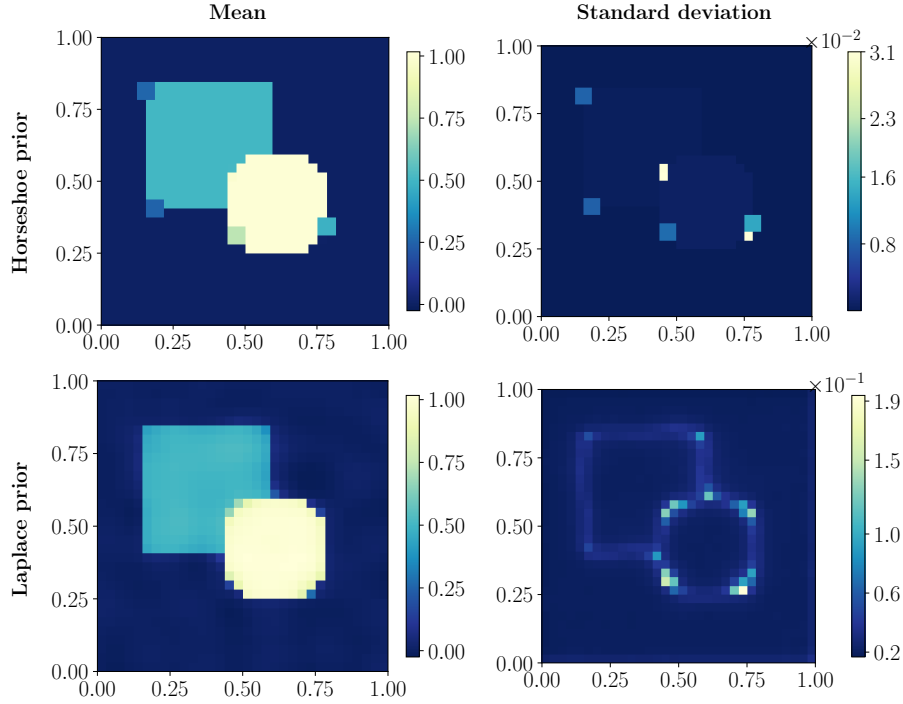


Figure 11: Comparison of posterior statistics for the target parameter  $\mathbf{x}$  using the horseshoe prior (first row) and the Laplace prior (second row). Posterior mean and standard deviation.

The discretized CT reconstruction problem takes the form given in (1), where  $\mathbf{x}^{\text{true}} \in \mathbb{R}^d$  is the vector of attenuation coefficients, which corresponds to an  $N \times N$  matrix with  $d = N^2$ , and  $\mathbf{y} \in \mathbb{R}^m$  denotes the measurement data. The system matrix  $\mathbf{A} \in \mathbb{R}^{m \times d}$  is from the discretized Radon transform, and the number of rows  $m$  is the product of the number of detector elements  $p$  and the number of projection angles  $q$  (see, e.g., [50, Ch.9] for the details). In this context,  $\mathbf{y}$  is referred to as projection data that is typically expressed as a 2D array known as sinogram.




Figure 12: Comparison between a few posterior samples generated from the horseshoe prior (first row) and Laplace prior (second row).

In particular, we consider the reconstruction of a “grains” phantom generated with the `phantomgallery` function from [51]. The projection geometry and data acquisition process comes from a *fan beam* configuration [50, p.53], whose parameters are defined as follows (in arbitrary units): the distance from the X-ray source to the origin is  $3N$ , the distance from the origin to the detector is  $N$ , and the number of detector elements is  $p = N$ . The discretized domain size is set to  $N = 64$  (thus  $d = 4096$ ), and we use  $q = 32$  equidistantly distributed projection angles in  $[0^\circ, 360^\circ]$ . Furthermore, the noise standard deviation is set as  $\sigma_{\text{obs}} = 0.01(\|\mathbf{Ax}^{\text{true}}\|_2 / \sqrt{m})$  with  $m = pq = 2048$ . Figure 13 shows the true image together with the sinogram data.




Figure 13: Grains phantoms. Left: underlying true image. Right: noisy sinogram data.

For the hyperparameters  $\sigma_{\text{obs}}$  and  $\tau$  we obtain a similar behavior to the previous examples (hence we omit the plot). The posterior mean and standard deviation of  $\sigma_{\text{obs}}$  and  $\tau$  are  $0.349 \pm 0.006$  and  $4.94 \times 10^{-6} \pm 3.84 \times 10^{-8}$ , respectively. The IACTs for both hyperparameters are close to 1, and thus, we also obtain quasi-independent samples using the proposed Gibbs sampler in this example.

The posterior mean and standard deviation for the local weight parameter  $\mathbf{w}$  are shown in Figure 14. Similarly to the deblurring example, we see that  $\mathbf{w}^{(1)}$  captures the locations of vertical edges and  $\mathbf{w}^{(2)}$  finds the horizontal edges. In this case, it is more noticeable the influence of the assumed zero boundary conditions on the image; see the leftmost column in  $\mathbf{w}^{(1)}$  and the topmost row in  $\mathbf{w}^{(2)}$ . We also observe that some of the edges are missed, especially where there is no significant differences in the attenuation coefficients of the grains (see, e.g., bottom right area in the posterior means).

The main posterior statistics for the parameter  $\mathbf{x}$  are shown in Figure 15. We plot the posterior mean and standard deviation, and compare them with the results from using the Laplace Markov random field prior. Again, we show the posterior mean with the same intensity range for the sake of comparison, but no negative values are obtained in the horseshoe case. Using the posterior mean as point estimate, the relative reconstruction error based on the mean are 0.045 for the horseshoe prior and 0.052 for the Laplace prior. Similarly to the other examples, it is clear that the horseshoe prior allows us to obtain a




Figure 14: Posterior statistics for the local weights  $\mathbf{w}^{(1)}$  (first row) and  $\mathbf{w}^{(2)}$  (second row). Posterior mean (left) and standard deviation (right).

sharper reconstructions with smaller posterior uncertainty compared to the Laplace prior. Despite some of the grain features are missed in the reconstruction, we observe that the standard deviation values reflect higher uncertainty at those particular missing locations. Moreover, note there are a few pixels at which the uncertainty is larger compared to the others; as in the deblurring example, these tend to be located at the corners of the grains.

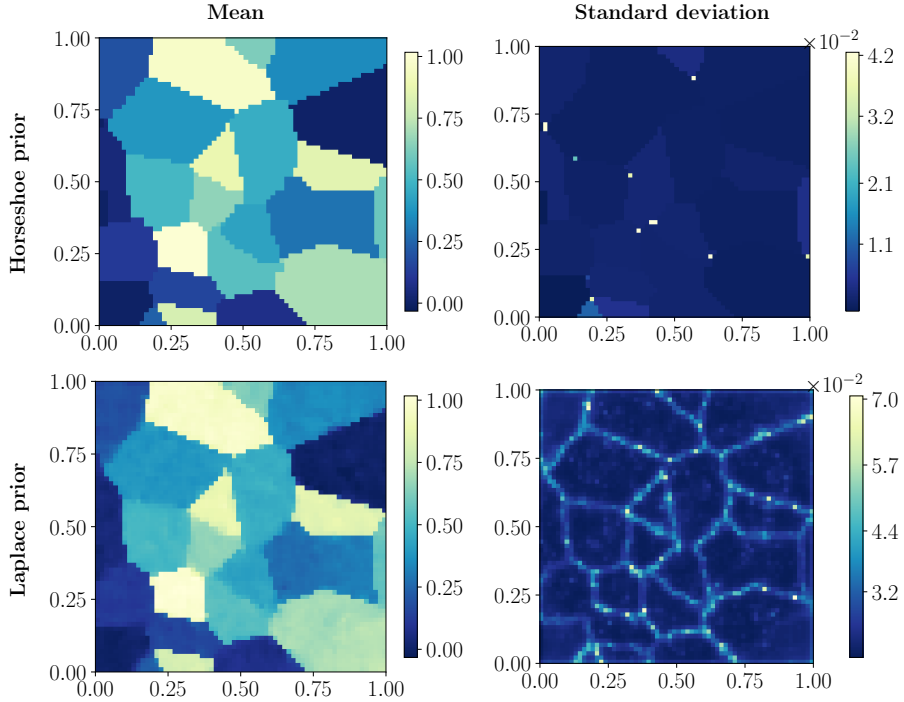


Figure 15: Comparison of posterior statistics for the target parameter  $\mathbf{x}$  using the horseshoe prior (first row) and the Laplace prior (second row). Posterior mean and standard deviation.

## 6. Summary and conclusions

We developed a computational framework for solving linear Bayesian inverse problems where preserving edges or sharp features of the solution is required. Our approach is based on the hierarchical horseshoe prior which combines a conditionally Gaussian prior on pairwise parameter increments with two half-Cauchy hyperpriors on global and local variance hyperparameters. The global hyperparameter reduces small increments towards zero, while the local hyperparameter allows some sufficiently large increments to escape shrinkage, thereby detecting the edges.

Due to the hierarchical structure and heavy-tailed distributions involved in the prior, sampling of the posterior distribution is cumbersome. Hence, we employ a scale mixture representation of the half-Cauchy hyperpriors to obtain a formulation with distributions of simpler form. This incorporates auxiliary parameters that allow improving convergence by embedding the target posterior into a larger model, as they assist in the reduction of dependence in the conditional specifications in the joint posterior (see, e.g., [52]). Thereafter, sampling of the posterior is performed via the Gibbs sampler where the conditional densities are derived in closed form by exploiting conjugate relations. In particular, the conditional density for the target parameter is high-dimensional Gaussian and we formulate the task of sampling as a least-squares problem that is solved efficiently via CGLS with preconditioner given by the prior precision matrix. We point out that despite our method is tailored for linear inverse problems, it can also be extended to nonlinear ones by sampling the conditional densities for the target parameter and the noise variance hyperparameter with suitable MCMC samplers.

The numerical experiments show that our approach based on the horseshoe prior enables the computation of sharp point estimates of the posterior, while reducing the posterior uncertainty. This compared with more traditional prior models based on Laplace Markov random fields. Furthermore, in one-dimensional applications the computational cost of the algorithm is significantly low since the preconditioner for CGLS is obtained as a byproduct of the formulation. However, in large-scale two-dimensional problems our approach becomes prohibitively slow since the preconditioner is expensive to compute.

Finally, some ideas to extend/improve our methodology are as follows:

- (i) Further work is needed to reduce the computational cost of sampling the Gaussian conditional in large-scale inverse problems. Specially, the analysis of its computational cost under mesh refinement. Methods from linear algebra such as inverse square root and matrix splitting approximations can be useful (see, e.g., [43, 23]), or MCMC methods that exploit low-dimensional structure [53], as well as, finding easy-to-compute preconditioners for CGLS.
- (ii) Application of the regularized horseshoe prior is a reasonable extension of our framework to make the posterior inference more robust (see, e.g., [41]). The idea is to control potential issues arising from the heavy-tails of shrinkage priors.
- (iii) The Bayesian inverse problem can be re-parameterized in terms of the increments (see, e.g., [20]), in this case, the prior precision matrix is equal to the weight matrix. Despite this transformation complicates the likelihood function, this can prove useful when computing the preconditioner required in pCGLS since the precision matrix is now diagonal. Moreover, in two-dimensional settings, the local weights associated to each coordinate direction can be combined to have the same size of the target parameter. The idea is to identify redundant edges by adding compatibility conditions (see also [20]).
- (iv) While it has been discussed that the horseshoe prior is typically resistant to problems originating from weak identifiability (see, e.g., [32]), theoretical and computational studies are required in connection with potential multi-modality of the posterior under the horseshoe prior, as well as, posterior concentration properties of the hierarchical horseshoe prior (see, e.g., [54]).
- (v) The number of hyperparameters can be reduced using empirical Bayes or maximum marginal likelihood methods. The idea is to fix the hyperparameter to the value that maximizes the model evidence. This strategy is used for example in sparse Bayesian learning [55].

## Appendix A. Definitions and scale mixture proof

We use this appendix to recall the definition of the inverse gamma and half-Student's  $t$  densities, and to show a simple derivation of the scale mixture representation given in (15).



First, the probability densities for the inverse gamma and half-Student's t random variables (discussed in section 3.2) are respectively defined as

$$\text{IG}(\alpha, \beta) \implies \pi(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{-\alpha-1} \exp\left(-\frac{\beta}{x}\right), \quad \text{and} \quad (\text{A.1a})$$

$$t^+(\nu, 0, c) \implies \pi(x) = \begin{cases} \frac{2\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\pi\nu}c} \left(1 + \frac{1}{\nu} \frac{x^2}{c^2}\right)^{-\frac{\nu+1}{2}}, & \text{if } x \geq 0 \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A.1b})$$

Second, the scale mixture (15) can be derived by marginalization of the associated hierarchical prior, that is

$$\pi(a^2) = \int_0^\infty \pi(a^2 | b) \pi(b) db \propto (a^2)^{-\frac{\nu}{2}-1} \int_0^\infty b^{-\frac{\nu+1}{2}-1} \exp\left(-\frac{1}{b} \left(\frac{\nu}{a^2} + \frac{1}{c^2}\right)\right) db \quad (\text{A.2a})$$

$$\propto (a^2)^{-\frac{\nu}{2}-1} \left(\frac{\nu}{a^2} + \frac{1}{c^2}\right)^{-\frac{\nu+1}{2}} \propto (a^2)^{-\frac{1}{2}} \left(1 + \frac{1}{\nu} \frac{a^2}{c^2}\right)^{-\frac{\nu+1}{2}}, \quad (\text{A.2b})$$

where we use the fact that the integrand in (A.2a) is an unnormalized inverse gamma density with shape and scale parameters  $(\nu + 1)/2$  and  $\nu/a^2 + 1/c^2$ , respectively. This allows analytical computation of the marginalization procedure. Finally, to obtain the density for  $A$  from that of  $A^2$ , we can use a probabilistic transformation. This results in the half-Student's t-distribution with degrees of freedom  $\nu$  and scale parameter  $c$ , defined in (15).

## References

- [1] P. C. Hansen, *Discrete Inverse Problems: Insight and Algorithms*, Society for Industrial and Applied Mathematics (SIAM), 2010.
- [2] J. Kaipio, E. Somersalo, *Statistical and Computational Inverse Problems*, Springer, 2005.
- [3] D. Gamerman, H. F. Lopes, *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*, 2 Edition., Chapman and Hall/CRC, 2006.
- [4] A. B. Owen, *Monte Carlo Theory, Methods and Examples*, [statweb.stanford.edu/~owen/mc/](http://statweb.stanford.edu/~owen/mc/), 2019.
- [5] R. Ranganath, S. Gerrish, D. M. Blei, Black box variational inference, in: *17th International Conference on Artificial Intelligence and Statistics*, 2014, pp. 1–9.
- [6] T. A. El Moselhy, Y. M. Marzouk, Bayesian inference with optimal maps, *Journal of Computational Physics* 231 (2012) 7815–7850.
- [7] C. Schillings, B. Sprungk, P. Wacker, On the convergence of the Laplace approximation and noise-level-robustness of Laplace-based Monte Carlo methods for Bayesian inverse problems, *Numerische Mathematik* 145 (2020) 915–971.
- [8] F. Uribe, J. M. Bardsley, Y. Dong, P. C. Hansen, N. A. B. Riis, A hybrid Gibbs sampler for edge-preserving tomographic reconstruction with uncertain view angles, *SIAM/ASA Journal on Uncertainty Quantification*. 10 (2022) 1293–1320.
- [9] A. M. Bruckstein, D. L. Donoho, M. Elad, From sparse solutions of systems of equations to sparse modeling of signals and images, *SIAM Review* 51 (2009) 34–81.
- [10] M. Lassas, S. Siltanen, Can one use total variation prior for edge-preserving Bayesian inversion?, *Inverse Problems* 20 (2004) 1537–1563.
- [11] J. M. Bardsley, Laplace-distributed increments, the Laplace prior, and edge-preserving regularization, *Journal of Inverse and Ill-Posed Problems* 20 (2012) 271–285.
- [12] J. Suuronen, N. K. Chada, L. Roininen, Cauchy Markov random field priors for Bayesian inversion, *Statistics and Computing* 32 (2022) 1573–1375.
- [13] M. Dashti, S. Harris, A. Stuart, Besov priors for Bayesian inverse problems, *Inverse Problems & Imaging* 6 (2012) 183–200.
- [14] M. Lassas, E. Saksman, S. Siltanen, Discretization-invariant Bayesian inversion and Besov space priors, *Inverse Problems & Imaging* 3 (2009) 87–122.
- [15] B. Hosseini, Well-posed Bayesian inverse problems with infinitely divisible and heavy-tailed prior measures, *SIAM/ASA Journal on Uncertainty Quantification* 5 (2017) 1024–1060.
- [16] M. M. Dunlop, M. A. Iglesias, A. M. Stuart, Hierarchical Bayesian level set inversion, *Statistics and Computing* 27 (2017) 1555–1584.
- [17] U. S. Kamilov, H. Mansour, B. Wohlberg, A plug-and-play priors approach for solving nonlinear imaging inverse problems, *IEEE Signal Processing Letters* 24 (2017) 1872–1876.
- [18] C. Li, M. Dunlop, G. Stadler, Bayesian neural network priors for edge-preserving inversion, *Inverse Problems and Imaging* 0 (2022) 1–26.
- [19] D. Calvetti, E. Somersalo, A. Strang, Hierarchical Bayesian models and sparsity:  $\ell_2$ -magic, *Inverse Problems* 35 (2019) 035003.
- [20] D. Calvetti, M. Pragliola, E. Somersalo, Sparsity promoting hybrid solvers for hierarchical Bayesian inverse problems, *SIAM Journal on Scientific Computing* 42 (2020) A3761–A3784.
- [21] N. G. Polson, V. Sokolov, Bayesian regularization: from Tikhonov to horseshoe, *WIREs Computational Statistics* 11 (2019) e1463.
- [22] S. van Erp, D. L. Oberski, J. Mulder, Shrinkage priors for Bayesian penalized regression, *Journal of Mathematical Psychology* 89 (2019) 31–50.

- [23] M. Vono, N. Dobigeon, P. Chainais, High-dimensional Gaussian sampling: A review and a unifying approach based on a stochastic proximal point algorithm, *SIAM Review* 64 (2022) 3–56.
- [24] T. Park, G. Casella, The Bayesian lasso, *Journal of the American Statistical Association* 103 (2008) 681–686.
- [25] Q. Li, N. Lin, The Bayesian elastic net, *Bayesian Analysis* 5 (2010) 151–170.
- [26] C. M. Carvalho, N. G. Polson, J. G. Scott, The horseshoe estimator for sparse signals, *Biometrika* 97 (2010) 465–480.
- [27] Å. Björck, *Numerical Methods for Least Squares Problems*, Society for Industrial and Applied Mathematics (SIAM), 1996.
- [28] A. Mohammad-Djafari, Bayesian approach with prior models which enforce sparsity in signal and image processing, *EURASIP Journal on Advances in Signal Processing* 52 (2012) 1–19.
- [29] M. Riis Andersen, A. Vehtari, O. Winther, L. K. Hansen, Bayesian inference for spatio-temporal spike-and-slab priors, *Journal of Machine Learning Research* 18 (2017) 1–58.
- [30] L. Tenorio, *An Introduction to Data Analysis and Uncertainty Quantification for Inverse Problems*, Society for Industrial and Applied Mathematics (SIAM), 2017.
- [31] C. M. Carvalho, N. G. Polson, J. G. Scott, Handling sparsity via the horseshoe, in: D. van Dyk, M. Welling (Eds.), *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*, volume 5, PMLR, 2009, pp. 73–80.
- [32] J. Piironen, A. Vehtari, Sparsity information and regularization in the horseshoe and other shrinkage priors, *Electronic Journal of Statistics* 11 (2017) 5018–5051.
- [33] H. Rue, L. Held, *Gaussian Markov Random Fields. Theory and Applications*, Chapman & Hall/CRC, 2005.
- [34] J. M. Bardsley, *Computational Uncertainty Quantification for Inverse Problems*, Society for Industrial and Applied Mathematics (SIAM), 2019.
- [35] E. Makalic, D. F. Schmidt, A simple sampler for the horseshoe estimator, *IEEE Signal Processing Letters* 23 (2016) 179–182.
- [36] M. P. Wand, J. T. Ormerod, S. A. Padoan, R. Frühwirth, Mean field variational Bayes for elaborate distributions, *Bayesian Analysis* 6 (2011) 847–900.
- [37] C. P. Robert, G. Casella, *Monte Carlo Statistical Methods*, 2 Edition., Springer, 2004.
- [38] D. Higdon, A primer on space-time modeling from a Bayesian perspective, in: B. Finkenstädt, L. Held, V. Isham (Eds.), *Statistical Methods for Spatio-temporal Systems*, Chapman & Hall/CRC, 2007, pp. 217–279.
- [39] M. Betancourt, M. Girolami, Hamiltonian Monte Carlo for hierarchical models, in: S. K. Upadhyay, U. Singh, D. K. Dey, A. Loganathan (Eds.), *Current Trends in Bayesian Methodology with Applications*, Chapman and Hall/CRC, 2015, pp. 79–101.
- [40] J. Johndrow, P. Orenstein, A. Bhattacharya, Scalable approximate MCMC algorithms for the horseshoe prior, *Journal of Machine Learning Research* 21 (2020) 1–61.
- [41] A. Nishimura, M. A. Suchard, Shrinkage with Shrunk Shoulders: Gibbs sampling shrinkage model posteriors with guaranteed convergence rates, *Bayesian Analysis* (2022) 1–24.
- [42] S. Geman, D. Geman, Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-6* (1984) 721–741.
- [43] E. Chow, Y. Saad, Preconditioned Krylov subspace methods for sampling multivariate Gaussian distributions, *SIAM Journal on Scientific Computing* 36 (2014) A588–A608.
- [44] D. Calvetti, F. Pitolli, E. Somersalo, B. Vantaggi, Bayes meets Krylov: Statistically inspired preconditioners for CGLS, *SIAM Review* 60 (2018) 429–461.
- [45] J. Besag, Spatial interaction and the statistical analysis of lattice systems, *Journal of the Royal Statistical Society. Series B (Methodological)* 36 (1974) 192–236.
- [46] G. O. Roberts, A. F. M. Smith, Simple conditions for the convergence of the Gibbs sampler and Metropolis-Hastings algorithms, *Stochastic Processes and their Applications* 49 (1994) 207–216.
- [47] L. Tierney, Markov chains for exploring posterior distributions, *The Annals of Statistics* 22 (1994) 1701–1762.
- [48] S. K. Bhattacharya, K. Khare, S. Pal, Geometric ergodicity of Gibbs samplers for the Horseshoe and its regularized variants, *Electronic Journal of Statistics* 16 (2022) 1–57.
- [49] P. C. Hansen, J. G. Nagy, D. P. O’Leary, *Deblurring Images: Matrices, Spectra and Filtering*, Society for Industrial and Applied Mathematics (SIAM), 2006.
- [50] P. C. Hansen, J. S. Jørgensen, W. R. B. Lionheart (Eds.), *Computed Tomography: Algorithms, Insight, and Just Enough Theory*, Society for Industrial and Applied Mathematics (SIAM), 2021.
- [51] P. C. Hansen, J. S. Jørgensen, *AIR Tools II: algebraic iterative reconstruction methods, improved implementation*, *Numerical Algorithms* 79 (2018) 107–137.
- [52] J. Besag, P. J. Green, Spatial statistics and Bayesian computation, *Journal of the Royal Statistical Society: Series B (Methodological)* 55 (1993) 25–37.
- [53] J. Martin, L. C. Wilcox, C. Burstedde, O. Ghattas, A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion, *SIAM Journal on Scientific Computing* 34 (2012) A1460–A1487.
- [54] S. L. van der Pas, B. J. K. Kleijn, A. W. van der Vaart, The horseshoe estimator: posterior concentration around nearly black vectors, *Electronic Journal of Statistics* 8 (2014) 2585–2618.
- [55] M. E. Tipping, Sparse Bayesian learning and the relevance vector machine, *Journal of Machine Learning Research* 1 (2001) 211–244.